Next Generation Stores of Value

Privacy, Proofs, Compute

A Conditional Monetary Thesis
for Verifiable Digital Infrastructure

Jason St George

> *A cypherpunk monetary stack: privacy by default, proofs by construction, and compute that only gets paid when anyone can verify it.*

Version 3.6

2026

Rights & License

© 2026 Jason St George.

This work (text and original figures) is licensed under **Creative Commons Attribution 4.0 International (CC BY 4.0)**. You may share and adapt the material for any purpose, including commercially, provided you give appropriate credit, provide a link to the license, and indicate if changes were made.

License: https://creativecommons.org/licenses/by/4.0/

Trademarks: "Next Generation Stores of Value," "VerifyPrice," "VerifyReach," "VerifySettle," "VerifyFlow," "Work Credits," and related marks are trademarks/service marks of the author and are not licensed under CC BY 4.0. No endorsement implied.

Third-party materials: All third-party trademarks, images, and excerpts are the property of their respective owners and are used under fair use or by permission where applicable; they are not covered by the CC license unless explicitly stated.

How to cite:

> Jason St George. *Next Generation Stores of Value: Privacy, Proofs, Compute*. Version 3.6. Licensed CC BY 4.0.

Abstract

Gold works when the network does not; Bitcoin works when the state does not. Neither is priced for the state this thesis is about: administrative repression under intact infrastructure, where the courts sit, payments clear, the lights stay on, and guarantees are honoured selectively at a price set by someone else. That band is narrow by construction, and the narrowness is the specification rather than an apology for it. Under sustained administrative and financial repression, a bearer asset may earn monetary premium if holding it preserves private settlement, portable proof, and access to verified compute after ordinary legal, custodial, and platform substitutes weaken. That premium is conditional: the full service path must remain stress-deliverable through independently substitutable physical routes; demand must not bypass the asset; a persistent, self-custodied, loss-bearing holder constituency must absorb residual financial risk; and the duration-bearing credit required to build the infrastructure must remain explicitly separate from the monetary object. Privacy, Proofs, and Compute may share that asset, but whether they should is an empirical architectural question rather than a premise.

The protocol and its price are different machines. The protocol machine supplies the triad and routes its demand through fees, burns, collateral, and scheduled issuance. The market machine wraps the resulting asset into custody claims, funds, derivatives, leverage, and allocation rules. Either machine can move without the other. Price is something to explain, not evidence to cite.

Fees, burns, and collateral establish a competitively priced cash-flow claim and a balance-sheet lockup—both of which a discounted-cash-flow valuation reproduces and neither of which is a monetary premium. The monetary claim rests on three mechanisms that lie outside such a valuation. **Pledgeability without due diligence**: cheap universal verification can make a position information-insensitive by *symmetric knowledge* — checking is free and public, so looking confers no edge on anyone — rather than by the symmetric ignorance every prior money-like asset has relied on. **Denomination**: the coordination function a service protocol can reach because the goods are produced by the system that issues the unit. And a state-contingent **holder-side service flow**: the bearer's ability to transact, prove, hold, and exit when the substitutes for proofs, courts, and custodians have stopped working. Against its strongest alternative — holding Bitcoin and converting on the day — that service reduces to a balance that stays spendable when acquiring one is no longer feasible, and an access experiment tests whether it does. Each mechanism is standard in its own literature. This thesis argues all three, sizes none of them, and publishes an instrument and a red line for each.

The competitor to an open verifiable stack is accordingly not fiat but a **competent closed sovereign stack**—state-integrated energy, industry, compute, payments, and identity—which may win on cost, uptime, and build speed while inverting who ends up sovereign. Two such stacks are now under construction rather than one. The marginal buyer of cheap verification is increasingly a machine delegating to another machine across trust perimeters, and the thesis measures whether a human still stands behind what that buyer holds. Duration-neutrality of the monetary object is a repression-resistance requirement, not a theory of capital formation: reserve collateral and the construction interval of plants are different functions, and proofs must not become the bond. The thesis publishes **eighteen red lines**: named, measurable conditions any one of which would retire it. Section 8.1 states the argument as a chain of nine links and names the links it rates weakest.

Extended Abstract

Post-Bretton Woods money is increasingly backed by *compliance* and *regulatory enforcement* rather than reserves or convertibility. Debt stocks that cannot be honored in real terms make financial repression and balance-sheet capture arithmetically attractive. The dependency now runs in both directions: fiscal capacity itself comes to rest on the market value of the asset complexes the state regulates, and the marginal demand for the reserve asset is increasingly commissioned by benchmark rules and delivered by leverage. Synthetic media and platform moderation undermine the assumption that "seeing is believing." AI concentrates cognition and compute behind hyperscaler chokepoints. Hardware, networks, identity, and custody systems are progressively integrated into state and platform surveillance infrastructure. In that environment, stores of value that depend on soft guarantees—custodians, editorial gatekeepers, platform labels, regulated wrappers, or vendor invoices—are brittle.

This thesis tests a narrower claim: **Privacy**, **Proofs**, and **Compute** can support an agency-preserving bearer asset that may earn monetary premium under repression, but only where the full stack delivers usable service under stress and a durable native holder base warehouses the asset's risk. Fees, burns, collateral, disciplined issuance, and public telemetry establish value *accrual*: a competitively priced cash-flow claim and a balance-sheet lockup, both of which a discounted-cash-flow valuation reproduces and neither of which is a monetary premium. The monetary claim rests on the three mechanisms the Abstract names—pledgeability without due diligence, denomination, and a state-contingent holder-side service flow—each argued, none sized, and each given an instrument and a red line. The holder-side flow meets its strongest alternative first: hold Bitcoin and convert on the day. Against it, holding the base asset buys one thing, a balance that stays spendable where acquiring one is unavailable, too slow, or incompatible with the transaction's privacy and control terms while native settlement and the service path still work. It buys no capacity, price, or priority; those are sold through Work Credits, prepaid claims on service reserved with named providers. An access experiment tests whether the advantage exists against the best Bitcoin-funded strategy, provider-side conversion included, and whether it supports anything beyond operating inventory (Section 13.1). The thesis makes the mechanisms' physical, market, and holder conditions measurable so the question can be settled rather than asserted. Its threat model assumes intentional repression, not benevolence, including the softer **administrative repression** of custody defaults, compliance rules, and platform terms of service. Against that adversary it sets a **seven-layer architecture**. Layer 0 grounds it in verifiable machines, energy, and physical capacity (Facility Capacity Receipts); the layers above run through communications, software distribution, identity, truth

& work, value & settlement, and a telemetry/governance layer that measures drift rather than denying it.

At the economic core, the thesis formalizes **verification asymmetry** and defines **VerifyPrice** (the time-and-cost vector for independently checking a claim) as the hinge that turns proofs and verified FLOPs into commodities rather than platform IOUs. It proposes a **Create/Compute Prove Settle Verify** loop, a modular stack of **twelve primitives**, and **four reference applications**: private treasury & payroll, media provenance, verified inference, and proof/-compute procurement. An operator/investor telemetry program, with extensions such as **VerifyReach**, **VerifySettle**, and the **Wrapper Dominance Ratio**, makes neutrality, repression-resilience, and co-option risk falsifiable.

Even when these conditions hold at the protocol layer, observed asset prices can be dominated by external financial wrappers, price-insensitive allocations, mechanical rebalancing, and dealer hedging. This thesis therefore distinguishes **native value capture** from **market price capture** and introduces **VerifyFlow**, a public market-realization telemetry program measuring wrapper exposure, creations and redemptions, holder flow elasticity, leverage, concentration, dealer dependence, and the divergence between financial exposure and protocol-native use. Price is not treated as proof of monetary adoption. Recurring native fees, burns, collateral, settlement, and verified-work demand under healthy SLOs are what make value accrue to a triad asset at all, and are necessary before any monetary question can be asked of it. Wrapper-led appreciation without those conditions is financialization rather than monetary validation.

The buyer the thesis models is increasingly not a person. The marginal purchaser of cheap verification is a machine delegating to another machine, several hops from its principal, at a speed that leaves no room for due diligence, reputation, or courts. Whether that demand ever reaches an open stack depends on whether delegation crosses trust perimeters or stays inside one operator's fleet; whether its holdings count as anyone's depends on whether a human can still revoke, cap, and re-key the agent (Chapter 7). This gives the first link of the chain a mechanism and the holder link a caveat. It supplies no precedent.

Beneath both loops sits a physical dependency the thesis makes explicit rather than assuming away. Gross capacity is not deliverable service: the relevant quantity is the surviving flow through energy or fuel $\rightarrow$ firm electricity $\rightarrow$ transformer and switchgear $\rightarrow$ cooling and water $\rightarrow$ hardware $\rightarrow$ network $\rightarrow$ workload $\rightarrow$ proof $\rightarrow$ verification $\rightarrow$ settlement $\rightarrow$ usable service. The thesis therefore reports **Delivered Verified Capacity** beside **Sovereign Optionality**: the first measures scenario-specific surviving service flow and its binding constraints, while the second measures pathway diversity. Service claims are bounded by stress-adjusted deliverable capacity, not nameplate hardware or a benign-state average. Because verification affordability is denominated in real resources on reference hardware, it is exogenous to token price but *not* exogenous to the conditions under which power and unprivileged hardware can be obtained.

The hinge is externally triggerable.

The **Topological Scarcity Lemma** follows: a nominal stock of resources, capital, or capacity does not secure a monetary service; the relevant quantity is the stress-deliverable flow through the narrowest non-substitutable edge connecting input to holder utility. The associated **Pressure–Capacity Corridor** writes the potential monetary service as $\Pi_{\text{monetary}}(R) \propto \Delta(R) \times A_{\text{full-stack}}(R) \times H(R)$, where differential demand may rise with regime pressure, full-stack availability may fall, and $H(R)$ is the holder's practical ability to self-custody, transact, prove, and exit. The likely shape is an inverted U: little differential value under benign institutions, the strongest case while substitutes weaken but the open stack still works, and declining service when extreme pressure disables the stack or the holder.

The competitor to an open verifiable stack is accordingly not fiat alone but a **competent closed sovereign stack**—state-integrated energy, industry, compute, payments, and identity—which may win on cost, uptime, and build speed while inverting who ends up sovereign. Architecture is likewise not pre-selected: the thesis compares one native monetary asset, a neutral reserve plus service-specific credits, and shared settlement plus modular domain collateral. Capability was never the claim; the properties are. In total the thesis publishes **eighteen red lines**: named, measurable conditions any one of which would retire it.

Gold, silver, Bitcoin, and physical hard assets remain load-bearing bridge assets while this stack is under construction; the triad becomes a candidate destination only when VerifyPrice, VerifyReach, VerifySettle, Layer 0 capacity telemetry, and value-capture metrics prove resilience under adversarial conditions. This is not a call to abandon hard assets prematurely, nor a claim that new money is inevitable. Bitcoin's work function is conceded to be monetarily superior to proof-of-useful-work, and for a reason that does not go away: nobody outside the system buys a hash, so nobody can subsidize, mandate, or withdraw demand for it, and a useful-work function necessarily introduces a buyer who can be. The claim made here is not that a better money is available, but that portable attestation and verified compute are separate goods with separate demand, and that whether they carry monetary premium or merely price as services is open. The result is not a single chain or guaranteed new money, but a **research and engineering agenda**: a "Bell Labs" for proof-of-useful-work and lawful privacy, testing whether Privacy, Proofs, and Compute can earn a store-of-value premium under measurable conditions. The conditions are that verification remains cheap, privacy settlement remains usable, useful-work markets avoid capture, physical infrastructure remains legible, and protocol design converts recurring triad demand into scarce, non-bypassable, agency-preserving asset value.

Executive Memo: Why AfterFiat Exists

The fiat problem is no longer only inflation. It is the merger of debt arithmetic, compliance-backed money, platform-mediated truth, AI enclosure, and physical compute constraints.

Sovereigns with debt loads that cannot be honored in real terms will prefer repression to default. Platforms facing synthetic media and regulatory pressure will prefer managed truth to open verification. AI hyperscalers will prefer renting cognition to distributing it. Custodians and wrappers will prefer price exposure without native usage. Citizens below the participation line will be offered management rather than agency.

This thesis is the counter-architecture: private settlement instead of surveilled accounts, portable proofs instead of platform labels, verified compute instead of vendor promises, and public telemetry instead of trust. Capturing indispensable demand, and not being bypassed by fiat, stablecoins, cloud contracts, custodial wrappers, or administrative convenience, is what makes value *accrue* to the asset. It does not by itself make the asset money: a fee stream and a collateral lockup are quantities a discounted-cash-flow valuation reproduces. The monetary claim rests on three mechanisms a discounted-cash-flow valuation does not reach: whether a counterparty will accept the asset without investigating it, whether anyone will reckon in it rather than merely settle in it, and what the bearer can still do when the substitutes for proofs, courts, and custodians have stopped working. This thesis argues all three and sizes none of them.

The thesis is conditional, and two concessions belong at the front rather than in an objections chapter. Bitcoin's work function is monetarily superior to proof-of-useful-work, because a puzzle nobody outside the system buys has no demand curve anyone can move. And gross native demand is not a monetary anchor: just-in-time fee purchases, burns, operator inventory, collateral requirements, and wrapper exposure do not establish a persistent, self-custodied, loss-bearing constituency willing to add through stress. Gold, Bitcoin, and hard assets are the bridge. Privacy, Proofs, and Compute supply different goods; whether one native asset, a neutral reserve plus service credits, or shared settlement with modular collateral best serves them remains an empirical question.

The buyer is a delegation.

The marginal buyer of cheap verification is increasingly a machine acting for a principal several hops away: an agent delegating to another agent, at a speed that leaves no room for due diligence, reputation, or courts. Whether that demand ever reaches an open stack depends on whether delegation crosses trust perimeters or stays inside one operator's fleet, and that is set by the same verification cost the rest of the argument hinges on (Chapter 7). The structure of

this buyer is measured, not assumed: who operates it decides how coercible it is, and whether a human can still revoke it decides whether its holdings count as anyone's. It supplies no precedent and repairs no concession.

The protocol and its price are different machines.

The protocol machine supplies Privacy, Proofs, and Compute and routes their demand through fees, burns, collateral, and constrained issuance. The market machine wraps the resulting asset into custody claims, funds, derivatives, leverage, and systematic allocation rules.

The first machine determines whether the asset has monetary substance. The second determines how that substance—or the story about it—is represented in price.

Either machine can move without the other. A protocol can improve while its price falls during wrapper redemptions. Its price can rise while native use stagnates because an ETF, treasury company, or leveraged product directs large exposure demand toward a thin market.

AfterFiat therefore requires two forms of falsifiability: **stack telemetry** for monetary function, and **flow telemetry** for market realization.

A further unbundling belongs at the front rather than only in an objections chapter. Duration-neutrality of the monetary object is a repression-resistance requirement. It is not a theory of capital formation. Reserve collateral and long-lived project finance are different functions; Treasuries currently fuse them; proofs can audit the second and must not become it (Section 3.3, Section 34.5.1).

Both machines are bolted to a grid.

Cheap verification is the hinge of the entire argument, and it is priced in electricity and reference hardware. That makes it exogenous to token price but not to policy: a state that finds it awkward to ban verification can simply decline to energize it, and curtailment, rate discrimination, and interconnection queues survive judicial review comfortably. The exposures differ—verification is threatened by hardware obtainability, proving by bulk power—and conflating them produces overclaims in both directions. The thesis measures the substrate as **sovereign optionality**, prices its fragility into issuance, and retires itself if verification affordability becomes a sovereign policy variable rather than a market outcome.

The real competitor is competence, not fiat.

A state-integrated closed stack—sovereign energy, industry, compute, payments, and identity—is the same trust-minimization instinct applied to matter, and it may deliver better cost and uptime than anything open. The thesis does not contest that. It argues that non-custodial settlement, default privacy, portable identity, checkable receipts, and practical exit are what make an asset a store of value under repression, and that a closed stack withholds precisely those, because withholding them is what makes it closed. Identical dashboards, opposite answers on whether a user can leave. The greater risk is not losing to that stack but quietly becoming it, which is why agency and exit are on the scoreboard and why the thesis names **eighteen red lines** rather than promising success.

Why instrument rather than forecast.

Every prior restructuring of how civilization handles information arrived as an externality of something cheaper, and nobody ordered its consequences. The printing press was adopted to copy books. Nobody adopted it to create the silent private reader [Ong 1982], standardized parts [Eisenstein 1979], or the nation-state; those arrived over centuries, as byproducts of a new habit at scale. The lesson generalizes, with the caveat that it is mechanism rather than measurement: a new medium is adopted for the function it makes cheaper, and it restructures whoever habituates to it, on a schedule nobody plans.

The stack tested here is being adopted to check work more cheaply. What cheap checking may do to the sensory basis of trust—to who can verify whom, and therefore to who must be trusted at all—is similarly unpredictable, and this thesis does not pretend otherwise. It responds by instrumenting the transition instead of forecasting it: red lines, public boards, and SLOs exogenous to token price, so that the consequences, good or bad, announce themselves in telemetry rather than in narrative (Appendix J records what each revision added). The frame is borrowed from the media-transition literature [McLuhan 1964]; the falsification conditions are not.

Contents

Abstract ii

Extended Abstract iv

Executive Memo: Why AfterFiat Exists vii

Reader's Map xxiii

Thesis in Plain Language xxviii

The Conditional Chain at a Glance xxxi

Key Definitions xxxiv

I Context & Claims 1

1 Introduction 3
- 1.1 Central Claims . . . 6
- 1.2 Premises . . . 9
- 1.3 Contributions . . . 10
- 1.4 End-to-End Vignette: The Loop in Action . . . 13

2 The Failure of Soft Guarantees 17
- 2.0.1 A Prior Transition: What the Press Actually Changed . . . 18

3 The World Forces New Monetary Primitives 20
- 3.1 Macro Playbook: Debt Repression Flight to Neutrality . . . 20
  - 3.1.1 Then and Now . . . 22
  - 3.1.2 The Treasury Market as Control Panel . . . 22
  - 3.1.3 Who Actually Buys the Paper . . . 24
- 3.2 The Collateral Loop . . . 29
  - 3.2.1 The Circuit . . . 29
  - 3.2.2 Why This Explains Balance-Sheet Repression . . . 30
  - 3.2.3 A Branch in the Playbook . . . 30
  - 3.2.4 Why AI Sits at the Center of the Loop . . . 31

3.3 Two Kinds of Duration . . . 32
3.4 Where the Triad Enters . . . 33
3.5 Social: The Web's Trust Default Has Flipped . . . 34
3.6 Political: The Participation Line Has Broken . . . 35

4 First Principles: What a SoV Must Survive 37

5 Threat Model 41
5.1 Scope and Assets . . . 41
5.2 Adversary Classes . . . 42
5.3 Threats by Loop Stage . . . 48
5.4 Mapping the Authoritarian Playbook . . . 49
5.5 Energy & Physical Interdiction . . . 51
5.6 Explicit Tensions and Boundary Conditions . . . 53

6 Layers of the Cypherpunk Stack 57
6.1 Modules in Outline: Applications and Primitives . . . 60

7 The Delegation Economy: Who Buys Verification 62
7.1 The Marginal Buyer . . . 62
7.2 The Delegation Hop . . . 63
7.3 Verification Cost Sets Delegation Topology . . . 64
7.4 The Buyer Bifurcates . . . 65
7.5 Agents as Holders . . . 66
7.6 Standing . . . 67
7.7 Where a Hop Touches the Stack . . . 68
7.8 What This Chapter Exports . . . 69

II The Triad and the Monetary Candidate 70

8 The Triad and the Monetary Candidate 72
8.1 The Conditional Chain, End to End . . . 73
8.2 Topological Scarcity and the Pressure–Capacity Corridor . . . 76
8.3 The Triad Coherence Test . . . 78
8.4 Monetary Objects and Value Capture . . . 78
8.5 Three Reference Designs . . . 82

9 Privacy as Private Money 85

10 Proofs as Attestation Money 88

11 Compute Through the "AI Money" Lens **90**
11.1 Compute Deflation and What Remains Scarce 90

12 Work Credits: Reserved Service Claims **93**
12.1 Definition: Work Receipts vs. Work Credits 94
12.1.1 The Work Credit Lifecycle 95
12.2 Energy Anchoring 96
12.2.1 Layer 0 Maturity and Economic Consequences 96
12.3 Robots, AI, and the Demand for Work Money 98
12.4 Why Work Credits Remain Service Claims 98
12.5 Monetary Role 99
12.6 Issuance: Tying Credits to Real Capacity 99
12.6.1 The Supply Balance Equation 100
12.7 Energy-Priced, Not Energy-Pegged 102
12.8 Why Work Credits Are Typed Service Claims, Not Presumed Money 103

13 Why Hold the Base Asset? **105**
13.1 Why Hold Rather Than Convert on the Day? 105
13.2 Economic Linkage: How Triad Demand Becomes Asset Value 108
13.2.1 Value Capture vs. Price Capture 113
13.3 Fee Incidence and the Level of the Fee 113
13.3.1 Incidence: A Fee Is a Wedge on Turnover, Not a Claim on Profit 114
13.3.2 The Base, and Why It Deflates 115
13.3.3 The Bound on the Sustainable Fee 116
13.3.4 Bypass Is a Magnitude, Not a Binary 117
13.3.5 What Determines Δ, and Why It Is Probably Thin 118
13.3.6 The Tail Is Bimodal, and the Two Modes Are Mutually Corrosive 119
13.3.7 Two Bypass Channels the Thesis Has Not Been Counting 122
13.3.8 The Fork Is an Equilibrium, Not a Competitor 123
13.3.9 What the Five Conditions Establish, and What They Do Not 125
13.3.10 Pledgeability Without Due Diligence 127
13.3.11 Where a Monetary Premium Would Actually Come From 135
13.3.12 Collateral: Unit-Elasticity, Float, and Wrong-Way Risk 139
13.4 From Utility Demand to Monetary Premium 142
13.5 Why Users Cannot Simply Bypass the Asset 143
13.5.1 The Incumbent Bypass Channel Is Not Hypothetical 144
13.6 Wrapper Dominance Risk 145
13.7 The Market Realization Plane 147
13.7.1 A Three-Stage Model 148

13.7.2 The Numeraire Is Not Fixed . . . 149
13.7.3 Denomination: The Function the Thesis Has Never Measured . . . 151

14 Service-Instrument Design Space 156
14.1 The "ZK Money" Analytical Lens . . . 156
14.2 The "AI Money" Analytical Lens . . . 157
14.3 Hybrid Instruments . . . 157

15 The ZK + AI Service Economy 159
15.1 Why Open Hardware Is a Precondition for the ZK-Economy . . . 159
15.2 How This Opens a ZK + AI Economy . . . 160

16 SoV Evaluation Framework 162
16.0.1 Comparison: Traditional Hard Assets vs. the Base-Asset Candidate . . . 165

III Infrastructure Layers 0–3 167

17 Layer 0: Verifiable Machines & Energy 169
17.1 Why Layer 0 Is a Monetary Question . . . 169
17.2 Design Goals and Non-Goals . . . 170
17.3 Hardware as Base Reality for Work Credits . . . 171
17.4 The Layer 0 Feasibility Ladder . . . 172
17.4.1 Policy Hooks: Economic Treatment by Grade . . . 173
17.4.2 Collateral Haircut by Modality and Grade . . . 174
17.4.3 Quantitative Thresholds (Reference Design) . . . 176
17.4.4 Pragmatic Starting Point (L0-A) . . . 176
17.4.5 The Path Forward . . . 177
17.4.6 Why This Prevents "Layer 0 Is Impossible, Therefore Thesis Fails" . . . 177
17.5 Concrete Components of Layer 0 . . . 177
17.5.1 Open Designs Where Possible . . . 177
17.5.2 Attested Randomness and Entropy . . . 178
17.5.3 Attestation Without Priesthood . . . 178
17.5.4 Lot Sampling and Destructive Audits . . . 178
17.5.5 How This Opens New Economies . . . 183
17.6 Verifiable Power: Energy as First-Class Input . . . 184
17.7 Facility Capacity Receipts: Beyond Energy . . . 185
17.8 Sovereign Optionality: Aggregating FCR Into One Exposure . . . 186
17.8.1 Adaptation of Scale . . . 186
17.8.2 Mapping to Existing FCR Fields . . . 187

17.8.3 Facility and Network Form . . . 187
17.8.4 Governance of the Weights . . . 188
17.8.5 What $\mathcal{O}_s$ Is For . . . 188
17.9 Delivered Verified Capacity: Surviving Service Flow . . . 189
17.10 Resilience Versus Affordability: The Layer 0 Cost Tension . . . 192
17.10.1 Two Different Exposures . . . 192
17.10.2 Disruption-Adjusted VerifyPrice . . . 193
17.10.3 The Decision Rule . . . 195
17.11 Operational Patterns: Profiles, Upgrades, and Failure Modes . . . 196
17.11.1 Hardware Profiles and Workload Binding . . . 196
17.11.2 Upgrades and Deprecations . . . 197
17.11.3 PQ and Cryptographic Agility . . . 198
17.12 Stress Tests for Layer 0 . . . 198
17.13 What Layer 0 Exports to Higher Layers . . . 198

18 Layer 1: Reachability **200**
18.1 Threat Model: What the Stack Must Survive . . . 200
18.2 Design Rules: Protocol Posture Under Pressure . . . 201
18.3 Mechanisms: What Needs to Ship . . . 201
18.4 How Layer 1 Anchors the Triad . . . 203
18.5 VerifyReach: Communications Telemetry . . . 204
18.5.1 VerifyReach Measurement Specification . . . 204

19 Layer 2: Distribution & Execution **207**
19.1 Threat Model & Objectives . . . 207
19.2 Design Rules . . . 208
19.3 Reference Distribution Architecture . . . 209
19.4 Key Management & Release Process . . . 209
19.5 Update Economics & Neutrality . . . 210
19.6 Fallback Playbooks . . . 210
19.7 Execution Integrity . . . 210
19.8 Lawful Privacy in Distribution . . . 211
19.9 Update Telemetry . . . 211

20 Layer 3: Identity & Claims **213**
20.1 Why Identity Must Decouple from Doxxing . . . 213
20.2 Design Goals . . . 213
20.2.1 Issuer Pluralism Without Issuer Anarchy . . . 214
20.3 The Identity Kernel . . . 215

20.4 Human Identity: Personhood and Compliance Without Dossiers 215
20.4.1 Sybil Resistance Menu 216
20.5 Machine Identity: Capability, Not Brand 216
20.6 Reputation: Priced Behavior, Not Personal Data 218
20.7 Patterns That Keep Identity Private and Usable 218
20.8 Failure Modes & Guardrails 219
20.9 How Identity Fits the Modular Stack 219
20.10 Identity/Reputation SLOs 220
20.11 Implementation Sketch 220
20.12 Part III Summary: Minimum Viable Stack 221

IV Truth, Work, and Settlement: Layers 4–5 223

21 From Infrastructure to Economics 225

22 Layer 4: Truth & Work 226
22.1 What Layer 4 Is (and Is not) 226
22.2 Verification Asymmetry Revisited 226
22.2.1 Verification Modalities, and Which One Carries the Monetary Claim . . 227
22.3 Canonical Workloads and Proof Types 230
22.3.1 Canonical Workload Definition Template 231
22.3.2 SLA Tiers: Latency and Assurance 233
22.3.3 Workload Registry 234
22.4 PoUW Design Patterns 234
22.4.1 The Usefulness–Unpredictability Dilemma 237
22.5 Proof Factories and PaL 238
22.6 VerifyPrice Observatory 239
22.6.1 VerifyPrice Measurement Specification 240
22.7 VerifyPrice in Practice 242
22.7.1 The Constitutional SLO Band 242
22.7.2 Protocol Affordability SLOs (Operational) 244
22.7.3 Token-Quoted VerifyPrice (Market Signal, Not Target) 244
22.8 Physical VerifyPrice: Auditing the Infrastructure Behind a Workload 245
22.9 Layer-4 Stress Tests 246
22.10 Minimum Viable Layer-4 Economy 247

23 Layer 5: Value & Settlement 249
23.1 Settlement as a First-Class Workload 249
23.2 Design Constraints for Privacy Rails 250

23.3 The Privacy Rails Kit (PRK) . . . 250
23.4 Settlement Safety: Refund and Bridge Invariants . . . 251
23.5 VerifySettle: Settlement Telemetry . . . 253
23.6 Layer-5 Stress Tests . . . 253
23.7 Minimum Viable Layer-5 Corridor . . . 255

24 The Modular Stack **257**
24.1 Primitive Catalog: The Twelve Pieces . . . 257
24.2 The Four Reference Applications at a Glance . . . 260
24.3 Reference Application: Private Treasury & Payroll . . . 262
24.4 Reference Application: Media Provenance . . . 262
24.5 Reference Application: Verified Inference . . . 263
24.6 Reference Application: Proof/Compute Procurement . . . 264
24.7 Part IV Summary . . . 265

V Governance, Telemetry & Neutrality **267**

25 Layer 6: Governance & Telemetry **269**
25.1 "No Dashboards, No Trust": Public SLOs as Constitution . . . 269
25.2 Control Surfaces: Parameters, Upgrades, Emergency Powers . . . 270
25.3 Bell-Labs-Style R&D vs. On-Chain Governance vs. Off-Chain Norms . . . 270
25.4 Governance as SLOs: The Five Invariants . . . 272
25.5 Constitutional Enforcement . . . 274
25.5.1 Machine-Enforced Constraints (Hard) . . . 274
25.5.2 Override Path (Slow, Expensive) . . . 275
25.5.3 Emergency Path (Narrow) . . . 276
25.6 Monetary Constitution . . . 277
25.6.1 Base-Asset Issuance Envelope . . . 277
25.6.2 Work Credit Issuance Rules . . . 278
25.6.3 Fee Routing . . . 279
25.6.4 Retirement Rule . . . 279
25.6.5 Issuance Weights . . . 280
25.6.6 Coverage Target . . . 281
25.6.7 Value Capture Loop . . . 282

26 Extended Telemetry **284**
26.1 Reference Verifier Classes . . . 284
26.2 Reference Verifier Implementations . . . 285
26.3 Cost Vector Definition . . . 285

26.4 Sampling and Anti-Cherry-Pick . . . 286
26.5 Data Availability and Anti-Memory-Hole . . . 286
26.6 Reproducibility Contract . . . 287
26.7 How to Measure: Receipts, Not Vibes . . . 287
26.8 VerifyFlow: Instrumenting the Market Realization Plane . . . 288
26.8.1 Wrapper-Level Vector . . . 289
26.8.2 Asset-Level Vector . . . 289
26.8.3 Holder Cohorts, Vintage, and Policy Concentration . . . 292
26.8.4 Capital Survival: AUM Is Not Investor Return . . . 292
26.8.5 Dealer Balance Sheets and Hedge Paths . . . 293
26.8.6 Convenience-Yield Telemetry: Pricing the Holder-Side Service Flow . . . 293
26.8.7 Buyer Quality and Price Provenance . . . 295
26.8.8 Native Monetary Buyer Map . . . 296
26.8.9 Principal Map . . . 297
26.9 The Public Boards . . . 299
26.10 Detecting Drift . . . 301
26.11 Making Neutrality and Repression-Resilience Falsifiable . . . 301
26.12 Policy-Attack Stress Harness . . . 302
26.12.1 Market-Realization Stress Tests . . . 302

27 Legal, Policy, and Jurisdictional Posture 304
27.1 Lawful Privacy as Protocol Design . . . 304
27.1.1 What Lawful Privacy Is Not . . . 304
27.1.2 Social Coercion Metrics: When "Optional" Becomes Mandatory . . . 305
27.2 Securities Law: Classification Is Transaction-Specific . . . 306
27.3 Defense-in-Depth Against Bans and Sanctions . . . 308
27.4 Labs, Foundations, and Neutral Router Commitments . . . 310

28 Operator & Investor Checklist 311
28.1 Triad Supply: Privacy, Proofs, Compute . . . 311
28.2 Value Capture & Anti-Bypass . . . 311
28.3 Institutional Holdability . . . 312
28.4 Holders and the Monetary Mechanisms . . . 313
28.5 Stack Health: Layers 0–6 . . . 313
28.6 Telemetry Honesty . . . 314
28.7 Red Flags and Failure Patterns . . . 314

VI Dynamics, Risk & Implementation 316

29 Adoption Curve & Ecosystem Dynamics 318
29.1 The Transition Window: Hard Assets as Bridge . . . 318
29.1.1 Portfolio Migration: From Bridge Assets to Triad Assets . . . 319
29.2 Macro Policy State Model . . . 320
29.3 Four Phases of Adoption . . . 322
29.3.1 Phase Gates and Failure Gates . . . 322
29.3.2 Adoption Curve To Expect . . . 325
29.3.3 Phase I: Cypherpunk-Led (Prove the Plumbing in Public) . . . 325
29.3.4 Phase II: Allocator-Led (Proofs and Privacy Become Budget Lines) . . . 327
29.3.5 Phase III: Composability-Led (The Stack Disappears into Infrastructure) 329
29.3.6 Phase IV: Policy-Led and Path-Dependent (Verification as Public Good) 330
29.3.7 The Negative Case: Useful Infrastructure Without Money . . . 332
29.4 Failure Gates . . . 332

30 Risk Analysis & Failure Modes 333
30.1 Technical Risks . . . 333
30.2 Economic Risks . . . 334
30.3 Political Risks . . . 334
30.4 Spec Drift, Vaporware, and Scoreboard Capture . . . 335
30.4.1 Scoreboard Capture in the Market Realization Plane . . . 335
30.5 Collateral Architecture Under Wrong-Way Risk . . . 336
30.6 Red Lines: When the SoV Thesis Fails . . . 336
30.6.1 Why the Hedge Bid Reaches the Base Asset and Not Work Credits . . . 351
30.7 Market Realization Warnings . . . 359

31 Implementation Sketches for Builders 361
31.1 Minimum Viable Stack (MVS) Sequencing . . . 361
31.2 Layer-0 Verifiable Machines . . . 362
31.3 Privacy Rails: Making Settlement Safe and Boring . . . 362
31.4 Proof Factories: Receipts as a Service . . . 363
31.5 Compute Consensus Pilots . . . 363
31.6 Developer Kit: Make “Import Proofs” the Default . . . 364

32 Kardashev Labs: Build and Measure 366
32.1 What Is Being Bought: Stack First, Coin Optional . . . 366
32.2 What Kardashev Labs Is . . . 367
32.3 Deliverables on a Dated Horizon . . . 368
32.4 Team Shape . . . 370

32.5 Planning Envelope . . . 370
32.6 Entity, Licensing, and What the Lab Owns . . . 371
32.7 Research Agenda: The Problems Named as Open . . . 372
32.8 How the Lab Fails . . . 373

33 The Closed Sovereign Stack 374
33.1 A State That Already Assumes Soft Guarantees Have Failed . . . 375
33.2 Trust Minimization Applied to Matter . . . 377
33.3 The Mirror: The Same Doctrine, Arriving in the West . . . 378
33.4 Open Weights Are Not Open Sovereignty . . . 380
33.5 The Inversion: Who Is Sovereign . . . 381
33.6 What This Validates . . . 383
33.7 What This Warns . . . 383
33.8 What This Chapter Does Not Claim . . . 384
33.9 What This Changes in the Stack . . . 385

34 Objections & Responses 389
34.1 Monetary Objections . . . 389
34.1.1 "Why not just Bitcoin?" . . . 389
34.1.2 "The buyer concession is repaired by agents." . . . 391
34.1.3 "This is just another chain / coin." . . . 392
34.1.4 "The triad does not belong in one token." . . . 393
34.1.5 "Utility does not imply money." . . . 394
34.1.6 "The asset can be bypassed." . . . 394
34.1.7 "Gold still beats this because gold needs no network." . . . 394
34.2 Technical and Infrastructure Objections . . . 395
34.2.1 "Users don't care about privacy or proofs." . . . 395
34.2.2 "PoUW will centralize / can't compete with hyperscalers." . . . 396
34.2.3 "Proofs do not prove truth." . . . 396
34.2.4 "They'll just shut off the internet / app stores." . . . 396
34.2.5 "This burns too much energy." . . . 397
34.2.6 "You will lose the competition for electrons." . . . 397
34.3 State, Political, and Governance Objections . . . 398
34.3.1 "Governments will never allow lawful privacy at scale." . . . 398
34.3.2 "Receipts will become surveillance tools." . . . 399
34.3.3 "Permanent receipts make obedience rational." . . . 399
34.3.4 "This just gives the state better compliance tools." . . . 400
34.3.5 "The closed sovereign stack simply wins." . . . 400
34.3.6 "A perfect ledger hardens an unequal starting state." . . . 401

34.3.7 "This is political preference dressed as analysis." . . . 402
34.3.8 "AI will be enclosed anyway." . . . 402
34.3.9 "Agents will be enclosed with the compute." . . . 403
34.3.10 "Governance will just re-centralize." . . . 403
34.4 Market-Structure and Co-option Objections . . . 404
34.4.1 "The asset will be co-opted, not bypassed." . . . 404
34.4.2 "This will just become another ETF." . . . 405
34.4.3 "The price went up, so the monetary thesis is working." . . . 405
34.4.4 "Wrappers increase access, so co-option is harmless." . . . 405
34.4.5 "Market mechanics are external, so protocol designers can ignore them." 406
34.5 Completeness Objections . . . 406
34.5.1 "Who finances the reactor?" . . . 406
34.5.2 The Open Duration Warehouse . . . 407

35 Why a Base Asset Behind These Could Become Money 409
35.1 As a Store of Value . . . 411
35.2 As Stack . . . 412
35.3 As Telemetry . . . 412
35.4 Trading One Base Reality for Another . . . 413

36 Conclusion: A Bell Labs for Privacy, Proofs, and Compute 414
36.1 A Quiet Arrival . . . 415
36.2 Two Loops, Not One . . . 415
36.3 The Question Under Test . . . 416
36.4 Equip, Don't Manage . . . 416

Sources 418

A Formal Model of Verification Asymmetry & VerifyPrice 443
A.1 Definition 1: VerifyPrice(W) and Verification Overhead . . . 443
A.1.1 Implementation Note: VerifyPrice Telemetry Schema (Sketch) . . . 445
A.2 Definition 2: Facility Energy Receipt (FER) . . . 446
A.3 Definition 3: PIDL Receipt (Proof Interface Definition Language) . . . 447
A.4 Definition 4: VerifyReach(N, R) . . . 448
A.5 Definition 5: VerifySettle(C) . . . 448
A.6 Definition 6: Repression Wedge . . . 449
A.7 Definition 7: Liquidation Effect . . . 449
A.8 Definition 8: Physical VerifyPrice . . . 450
A.9 Definition 9: VerifyFlow . . . 451

B Practical KPIs & Telemetry Templates **452**
B.1 Instrument Table . . . 452
B.2 Operational Checklists by Domain . . . 464

C The SDK (Proofs-as-a-Library) **467**

D Energy & Plant Architecture **471**

E Hardware Profiles **473**
E.1 Concept . . . 473
E.2 Naming . . . 473
E.3 Hardware Profile Specification . . . 474
E.4 Hardware Profiles in Receipts and Policies . . . 474
E.5 Example Profiles . . . 474
E.6 What a Lab Would Need to Specify . . . 475

F Communications Resilience Mechanisms **478**

G Glossary of Terms & Notation **480**

H Formal Model of Market Realization, Wrapper Flows, and Price Capture **493**
H.1 Flow Elasticity . . . 494
H.2 Leveraged Rebalancing and Flow-Induced Exposure . . . 495
H.3 Net Mechanical Gain . . . 496
H.4 The Recycling Boundary . . . 496
H.5 Wrapper Recycling Ratio . . . 497
H.6 Wrapper-Class Applicability . . . 497
H.7 Return-Decoupled Allocation Flow . . . 499
H.8 Market Impact and Liquidity . . . 499
H.9 Volatility Drag on Leveraged Wrappers . . . 500
H.10 Dealer Hedge Decomposition . . . 501
H.11 Measurement Contract . . . 501
H.12 Reference Stress-Harness Sketch . . . 502
H.13 Boundaries and Caveats . . . 502
H.14 Buyer Quality and Price Provenance . . . 502

I Scenario Analysis: The Collateral Loop Under Stress **504**
I.1 Correlated Loss: A Recession With Rising Long Yields . . . 504
I.2 The AI–Fiscal Collateral Loop . . . 506
I.3 The National-Champion Conversion Pathway . . . 506
I.4 Quarantined Claims . . . 507

I.5 Indicators: What Would Confirm or Weaken This 507

J What Changed Across Releases **511**

Reader's Map

This thesis is structured as a ladder of seven interlocking layers. Each layer has one or more "home" sections, and most layers also surface in the **Create/Compute Prove Settle Verify** loop and in the telemetry program.

Use the following table as your compass:

Layer	Role in the Stack	Primary Sections	Key Metrics / Concepts	Falsification Test
Layer 0 Verifiable Machines & Energy	Physical base reality: open silicon, sampled supply chains, and the power plants that keep provers and routers alive. *For pure proof correctness, sound cryptography should tolerate untrusted provers. Layer 0 becomes essential when claims include physical capture, energy use, machine identity, witness confidentiality, or useful-work fairness.*	Chapter 17 *Layer 0* (open hardware, profiles, lot sampling, FERs); Chapter 31 (L0 implementation sketches); Appendix D (Energy & Plant); Appendix E (Hardware Profiles)	Hardware profiles (HIDs); lot-sampling coverage; Facility Energy Receipts (FERs); Facility Capacity Receipts (FCRs); sovereign optionality $\mathcal{O}_s$; Physical VerifyPrice (infra audit); disruption-adjusted VerifyPrice; PUE/ERE/WUE; kWh/receipt; profile incidents	RL10 (physical opacity) or RL13 ($\mathcal{O}_s^{\text{net}}$ below threshold / rationing jurisdictions); hinge VerifyPrice becomes sovereign-policy-driven

Continued on next page...

Layer	Role in the Stack	Primary Sections	Key Metrics / Concepts	Falsification Test
Layer 1 Reachability: Communications & Transport	Keeps packets flowing under DPI, filtering, and shutdowns. *Without Layer 1, the loop cannot run and "public verification" becomes theoretical.*	Chapter 18 *Layer 1* (threats, design rules, mechanisms, VerifyReach); Chapter 31 (comms sketch); Appendix F (Comms Resilience)	**VerifyReach(N,R)**: p50/p95 time-to-first-connection; $succ_1/succ_2$ reachability; failure causes; share encrypted/obfuscated traffic; peer-set diversity (geo/ASN); swap success under network stress	>20% of sampled ASNs in uncensored regions unreachable for >7 days
Layer 2 Distribution & Execution: Software Supply & Runtime	Ensures honest clients and updates can be shipped even when app stores, CDNs, and DNS are weaponized. *Protects the code that speaks the protocol and runs PaL/PRK.*	Chapter 19 *Layer 2* (threats, design rules, multi-home architecture, update pipeline, key mgmt, fallback playbooks); Chapter 31 (distribution tooling); Appendix B (update KPIs)	Update health: p50/p95 metadata & artifact fetch times; success/rollback/signature-failure rates; mirror/CDN top-N share; geographic/ASN diversity; key-rotation & incident history	Single key compromise can push malicious updates

Continued on next page...

Layer	Role in the Stack	Primary Sections	Key Metrics / Concepts	Falsification Test
Layer 3 Identity & Claims: Humans & Machines Without Doxxing	Lets humans and machines prove capabilities and rights (age, residency, uniqueness, model ownership, device profile) without turning identity into a global dossier. *Reputation is receipts, not biographies.*	Chapter 20 *Layer 3* (identity kernel, human & machine identity, reputation, guardrails, how identity uses the loop); Chapter 27 (lawful-privacy requirements); Appendix C (PaL/PRK SDK)	Identity-predicate VERIFYPRICE (p50/p95 verify times for age/residency/set-membership proofs); unlinkability tests; reputation from PIDL receipts (SLA fulfillment vs slashes); hardware-anchored machine attestations	Identity requires centralized issuer
Layer 4 Truth & Work: Proof Systems and PoUW	Converts expensive work into cheap-to-check receipts. *This is where verification asymmetry lives and where* VERIFYPRICE *is defined.*	Conceptual: Chapter 11 *Compute Through the "AI Money" Lens*, Chapter 12 *Work Credits*; Structural: Chapter 21 *Infrastructure to Economics*; Chapter 22 *Layer 4* (canonical workloads, PoUW, proof factories, PaL); Chapter 24 (primitive catalog); Appendix A (formal VerifyPrice)	**VERIFYPRICE(W)** per workload: p50/p95 verify times/costs; failure rates; $r(W) = v(W)/p(W)$; prover concentration; SLA attainment (Bronze/Silver/-Gold); MatMul-PoUW economics; Work-Credit issuance	VerifyPrice p95 exceeds SLO for $\geq$3 months

Continued on next page...

Layer	Role in the Stack	Primary Sections	Key Metrics / Concepts	Falsification Test
Layer 5 Value & Settlement: Privacy Rails & Non-Custodial Flow	Moves value non-custodially and privately, with auditability via receipts. *This is where the service markets described by the legacy "Private Money," "AI Money," and "Proof/Attestation Money" lenses settle; their Work Credits and capacity instruments remain typed service claims. The separate base asset is the only conditional monetary candidate.*	Analytical lenses: Chapter 9 *Privacy as Private Money*, Chapter 11 *AI Money*, Chapter 10 *Proof/Attestation Money*; service claims: Chapter 12 (Work Credits); the monetary argument: Chapter 13 (Why Hold the Base Asset?); Stack: Chapter 23 *Layer 5* (settlement as workload, PRK, bridge-safety, corridors); Chapter 24 (treasury/payroll, media, inference, procurement flows); Chapter 31 (privacy toolkit); Appendix F (swaps, routing)	**VERIFYSETTLE(C)**: swap success $\geq$ target; **refund_safe(C)** — a bounded-time no-loss exit for every party, zero protocol-attributable losses (Section 23.4); p50/p95 time-to-finality; per-asset anonymity metrics; corridor LP/route concentration; non-custodial vs custodial flow share	RL2: any protocol-attributable refund failure on an admissible route, or overall refund_safe $<$ 99.5% with failures unclassified (repeated, no delist)

Continued on next page...

Layer	Role in the Stack	Primary Sections	Key Metrics / Concepts	Falsification Test
Layer 6 Governance & Telemetry	The immune system and constitution: keeps drift and capture visible and forces responses. *Turns "trustlessness" into dashboards and runbooks.*	Metrics: Chapter 22 (VerifyPrice observatory); Chapter 26 *Extended Telemetry*; Governance: Chapter 25 *Layer 6* (SLOs as constitution, control surfaces, Bell-Labs vs on-chain vs norms); Legal: Chapter 27; Ops: Chapter 28 (checklist & red flags); Dynamics: Chapter 29 and Chapter 30; Appendix A and Appendix B	VerifyPrice dashboards per workload; **VerifyReach** & **VerifySettle** boards; **VerifyFlow** (market realization); twelve public boards or maps (Section 26.9), including native-buyer, common-cause, and sponsor-dependency telemetry; decentralization metrics; FeeCoverage (retained fees); Wrapper Dominance; agency preservation	Receipt datasets unverifiable or controlled by single party; wrapper-led price without native use (interpret via VerifyFlow, not as sole red line)

Readers who want the version history—what each release from v1.2 onward added, and why—will find it in Appendix J. It is not needed to read the argument. Readers who want to know what would actually be built, by whom, on what budget and against which acceptance tests, should go to the investment addendum, Chapter 32, which maps the Phase I deliverables onto the instruments this map lists.

Thesis in Plain Language

The core claim: Under sustained administrative and financial repression, a bearer base asset may earn monetary premium if holding it preserves private settlement, portable proof, and access to verified compute after ordinary substitutes weaken. The full service path must remain stress-deliverable; demand must not bypass the asset; a persistent, self-custodied, loss-bearing holder constituency must absorb residual risk; and infrastructure credit must remain separate. Privacy, Proofs, and Compute may share the asset, but the thesis does not assume they should.

The accrual mechanism: Service value can accrue to the base asset through four channels:

1. fees paid in the native asset for proof generation, privacy settlement, and verified compute;
2. burns that permanently remove supply when capacity is consumed;
3. required collateral for provers, routers, and liquidity providers; and
4. issuance on a schedule fixed ex ante, not fiat decree.

The value-capture bridge: Utility demand alone does not create store-of-value premium. In the reference design, five channels route service value to the asset:

1. the asset is the **required fee medium** for core triad services;
2. a meaningful share of fees is **burned or retired**;
3. operators must **stake the asset as collateral**;
4. gross issuance follows a **schedule fixed ex ante**; and
5. users **cannot obtain equivalent service quality** through bypass channels that avoid the asset.

These are the reference design's channels, each necessary to its route and none necessary to value capture in general. If one fails in this design, the system may be useful infrastructure while the base asset fails even as a service-value claim. If all five hold, the result is a cash-flow claim and a collateral lockup, not yet money.

The monetary bridge adds what the accrual bridge cannot supply. Two are preconditions: stress-deliverable service, and a holder constituency that retains self-custodied balances, bears losses, and accumulates through relevant pressure rather than buying fees just in time or holding wrappers. Three are the mechanisms by which a premium could actually arise, each standard in its own literature and each unsized here (Section 13.3.10, Section 13.3.11, Section 13.7.3):

1. **pledgeability without due diligence** — cheap public verification removes the private information about unit quality that a lender would otherwise price, so a position can be accepted no questions asked; the prediction is lower cross-lender haircut *dispersion* than on opaque comparables, among lenders who underwrite independently, not lower haircut levels (Red

Line 17);

2. **denomination** — a service economy whose goods are priced, produced, and settled inside one accounting perimeter has a coordination path to being reckoned in, which a bare bearer asset lacks; measured as free-choice denomination share, never the mandated complement (Red Line 18);
3. **a state-contingent holder-side service flow** — the bearer's ability to transact, prove, hold, and exit when the substitutes for proofs, courts, and custodians have stopped working, which accrues by holding rather than by distribution; against Bitcoin converted on the day it reduces to a balance that stays spendable when acquiring one is no longer feasible, tested by an access experiment (Section 13.1) and measured as the convenience yield's increment in the pressure corridor (Red Line 16).

The falsifiable test: It runs one way. If VERIFYPRICE rises, VERIFYREACH falls, VERIFYSETTLE breaks, DVC collapses, the capturable wedge closes, the native holder base proves procyclical, the collateral–capacity spiral breaches Red Line 15, or any of the three mechanisms reads null on its instrument (Red Lines 16–18), or the access experiment finds that Bitcoin converted on the day does as well as a held balance, monetary candidacy fails. Healthy readings merely keep the candidate alive; no accumulation of green dashboards proves monetary premium.

What price does not tell you: None of the tests above is a price test. An asset's price can rise because ETFs, treasury companies, leveraged products, passive mandates, or dealer hedging direct large exposure demand at a thin market, while native fees, burns, settlement, and proof purchases stay flat—and it can fall while all of those improve. That is why the thesis adds a fourth measurement family, **VerifyFlow** (Section 26.8), to separate protocol-native demand from wrapper-led exposure. Price is something to explain, not evidence to cite.

Who is buying: Increasingly, software agents acting for people several steps away. An agent cannot investigate what it accepts at the speed it works, so it accepts most readily what anyone can check, which is why its demand is demand for verification. The thesis measures whether that demand crosses between operators, and whether a human can still revoke the agent that holds the asset.

The thesis examines the triad from three complementary angles, plus one plane around them:

- **Angle 1 Monetary:** whether a bearer base asset can earn a next-generation store-of-value premium by preserving agency through Privacy, Proofs, and Compute; the service claims themselves remain non-monetary.
- **Angle 2 Stack:** the seven-layer cypherpunk stack that actually supplies these capacities.
- **Angle 3 Telemetry & Governance:** VerifyPrice/Reach/Settle + ops as the thing that keeps "repression-resilient neutrality" falsifiable.
- **The Market Realization Plane:** the external financial machinery—wrappers, custody, lever-

age, dealers, allocation rules—that determines how, when, and with what volatility any of the above shows up in price (Section 13.7).

The Conditional Chain at a Glance

The thesis establishes its claim through a conditional chain rather than an assertion. The chain is defended link by link across six Parts, which makes it easy to lose. This page states it once; Section 8.1 walks it in full, says where each link is argued, and closes with Table 8.1, which rates the argument link by link.

> Utility demand → standardized work → receipts → stress-deliverable service → markets → fee flows → a scarce asset (only under non-bypassable accrual) → a persistent, self-custodied, loss-bearing, regime-responsive holder constituency → a possible store-of-value premium.

Every arrow is a conditional. None is asserted as inevitable, and the chain is only as strong as its weakest link. Each link therefore comes with the condition that would sever it:

	Link	Severed if
1	Utility demand exists and persists, through every cycle and especially under repression; machine delegation across trust perimeters is one named mechanism (Chapter 7)	the demand turns out to be cyclical discretionary spend rather than structural necessity, or delegation stays inside fleets that need no portable receipt
2	Demand standardizes into canonical work—workload registries, hardware profiles, SLA tiers	workloads stay bespoke, so no fungible unit forms and there is nothing to price
3	Work produces cheap-to-check receipts, with $r(W) = v(W)/p(W) \ll 1$	verification cost creeps toward production cost, and "anyone can verify" collapses into "trust the prover" (Red Line 1, the hinge)
4	Receipts represent stress-deliverable service, measured by scenario max-flow across the full stack, with binding constraints reported	nameplate capacity survives on paper while substitution latency or common-cause failure collapses usable service
5	Receipts enable markets: proof factories, neutral routers, SLA escrow, capacity procurement	routers, provers, or matching engines capture the market and rent replaces work (Red Line 3 territory)
6	Markets produce recurring fee flows, visible as fee coverage against the security budget	fee coverage collapses and workload demand proves speculative rather than budgeted
7	Fee flows accrue to a scarce asset—*only* under the five non-bypassable conditions of the Value Capture Lemma	users obtain equivalent triad capacity through hyperscalers, stablecoins, or custodial APIs without touching the asset (Red Line 6)
8	A persistent, self-custodied, loss-bearing, regime-responsive holder constituency anchors the asset	native demand is just-in-time, immediately sold, wrapper-led, leveraged, or procyclical, or the cohort is one operating policy at many addresses
9	A scarce, captured, anchored asset may earn a store-of-value premium, if it stays liquid, neutral, verifiable, legally holdable, and agency-preserving	any store-of-value requirement fails, a held balance does no better than Bitcoin converted on the day, or the telemetry showing it is captured (Red Line 4)

Where it is weakest. Six links are rated *Medium* or *Medium (conditional)*. Standardized work and proof/compute market formation are engineering-maturity risks. Stress-deliverable service is Medium because common-cause and substitution data are immature; non-bypassability, because the economics requires a magnitude rather than a binary; the holder anchor, because native transactions do not prove risk absorption; and the premium itself is conditional on the whole chain. The store-of-value premium derives from link 9 through three mechanisms

— pledgeability without due diligence, denomination, and the state-contingent holder-side flow — not from fee accrual (Section 13.3, Section 13.3.10, Section 13.3.11, Section 13.7.3); Red Lines 16–18 read on them. The full ratings are in Table 8.1.

What the chain does not claim. Completing all nine links would establish only that the base asset is a defensible monetary candidate. It would not establish when, or with what volatility, price reflects that (Corollary 13.1). Nor is the chain suspended in mathematics: the service links consume bulk power and hardware, and the holder link requires balance sheets capable of bearing loss. Both can fail while gross capacity or gross native demand still looks abundant.

Key Definitions

The core constructs that recur throughout this thesis are defined precisely here, before they are used. These are not metaphors; they are operationally specified primitives.

Definition: VerifyPrice(W) Specification Stub

For a canonical workload W, **VerifyPrice** is the public KPI vector:

$$\textsc{VerifyPrice}(W) \equiv \big(p_{50,t}(W),\ p_{95,t}(W),\ p_{50,c}(W),\ p_{95,c}(W),\ \mathrm{fail}(W)\big)$$

Where:

- $p_{50,t}(W)$, $p_{95,t}(W)$: median and 95th-percentile verification **time** (seconds)
- $p_{50,c}(W)$, $p_{95,c}(W)$: median and 95th-percentile verification **cost** (see cost vector below)
- $\mathrm{fail}(W)$: verification failure rate (fraction of attempts that fail or timeout)

This is the canonical five-field tuple; the cost fields derive from per-run receipts via the published price model of Section 22.6 (per-run cost, then percentiles — never recomputed from percentile inputs), and the physical-component tuples published alongside are views of this vector, not competing definitions (Appendix A).

Why this matters: VerifyPrice is the hinge that determines whether proofs and verified compute behave as commodities (publicly checkable) or as platform IOUs (trust someone's claim). If $r(W) = v(W)/p(W) \ll 1$, verification is cheap relative to production and markets can form; if $r(W) \to 1$, the market is back to "trust the prover."

Definition: Work Credits

A **Work Credit** is a typed, prepaid claim on reserved future service — privacy settlement, proof generation, or verified compute — owed by a specified provider or provider pool. Its issuance requires an admissible capacity commitment and an atomic reservation against outstanding obligations. Verified fulfillment extinguishes the claim and produces a receipt. The receipt is evidence of fulfillment and creates no further service obligation.

Lifecycle:

1. **Commit.** A named provider or constituted pool commits capacity for workload W at tier T over a stated delivery window, with Facility Energy Receipts and Facility

Capacity Receipts behind it.

2. **Sell and reserve, atomically.** The sale and the reservation are one step: no valid sold claim exists without its capacity reserved against every outstanding claim competing for the same capacity.
3. **Deliver and verify.** Delivery passes the registry acceptance criterion, with VerifyPrice(W,T) and the other SLOs within bounds.
4. **Extinguish and receipt.** Verified delivery retires the claim and issues the Work Receipt.

Cancellation, expiry, default, and refund follow the terms written in the claim. A failed delivery becomes neither an open-ended promise nor a claim on the base asset.

Uses: a credit can be redeemed for its service, transferred, or pledged as collateral where its tier permits (Section 22.10). Two neighbours are not Work Credits: a transferable record of historical work is a Work Receipt, which is evidence and not a claim, and advance proceeds that finance capacity not yet built are project finance, named and treated separately (Section 3.3). Work Credits carry *no* governance weight: the franchise is time-locked base asset only (Section 25.5).

Energy anchoring: A credit cannot be sold without an admissible capacity commitment, and it cannot be extinguished without delivering work that passes verification, so the cost of honoring one is bounded below by the energy and hardware that work requires. The difference from SHA-256 PoW is that this work has an *external buyer*—which is the point of the design and also its cost, since a buyer can be regulated, subsidized, or coerced in a way that a hash puzzle's absent buyer cannot (Section 34.1.1). Each credit references a Facility Energy Receipt (FER) chain; if the referenced plant drifts out of profile (PUE > 1.5, carbon intensity $>$ threshold, etc.), the credits reserved on it are flagged.

Deliverability bound: Outstanding claims are capped by stress-adjusted DVC for the specified workload and tier over the delivery horizon, not by gross energy, installed hardware, or benign-state proof throughput. DVC supports the reservation limit; it does not guarantee performance.

Non-debt property: Work Credits do not promise fixed coupons or redemption in fiat terms. Value floats with demand for triad capacity.

Failure mode: If VerifyPrice regresses materially, new sales halt until SLOs recover. Outstanding credits remain claims on their obligors under their stated terms and may trade at a discount, reflecting the network's degraded utility.

Definition: Instrument Hierarchy

The thesis distinguishes five categories of objects:

Evidence Objects (not scarce, not money):

- **PIDL Receipt**: Proof that a specific interaction occurred. Copyable, verifiable, not scarce.
- **Work Receipt**: Proof that a unit of work was completed. Evidence of past work; confers no future rights.
- **FER/FCR**: Evidence about energy, infrastructure, resilience, and the edge capacities used to compute DVC.

Capacity and Service Objects (typed claims, not presumed SoV):

- **Work Credit / WC-Voucher**: Prepaid or transferable access to specified triad capacity under workload, location, hardware, SLA, and delivery terms. May expire. Useful for service procurement and hedging, not presumed savings.

Duration-Bearing Credit:

- **Project Notes and Capacity Bonds**: Explicit credit claims with maturity, covenants, default states, and loss waterfalls. Proof-audited, but never transferred onto the base asset.

Derivatives and Operating Claims:

- **LP/Staking Shares**: Positions granting fee, slashing, corridor, or validator exposure. They are not monetary objects.

Conditional Monetary Candidate:

- **Base Asset**: The native fee and settlement unit. It is evaluated for monetary premium only if DVC, non-bypassability, holder quality, agency, and the remaining chain conditions pass.

Hierarchy rule: Throughout this thesis, "the asset" refers to the **base asset** unless otherwise specified. Evidence, service claims, duration-bearing credit, and derivatives do not inherit its conditional monetary candidacy.

Definition: Lawful Privacy

Lawful privacy is the design principle: *default privacy with optional, user-controlled disclosure*.

Concretely:

- **Default state:** Transactions, identities, and flows are encrypted and unlinkable without explicit consent.
- **Disclosure mechanisms:** Viewing keys, auditable receipts, and selective-disclosure proofs allow holders to prove specific facts (e.g., "I paid X to Y for purpose Z") without

exposing the full transaction graph.

- **No backdoors:** The protocol has no protocol-level master key, regulatory escrow, or "lawful intercept" API. The holder's own spending-key root is a master key for its own tree and is never disclosed; per-epoch keys and per-note proofs are the units of disclosure (Section 23.7). Disclosure is always at the holder's discretion.

Coercion boundary: Lawful privacy is a *technical* guarantee. It cannot prevent social or legal coercion to disclose viewing keys. What it guarantees is that (1) non-custodial routes exist, (2) disclosure cannot be forced at the protocol level, and (3) coercion surface is minimized by keeping data encrypted by default.

Quick Reference: Political Economy & Physical Capacity

These terms carry the thesis's political-economy and physical-capacity vocabulary. Full treatments appear at the cited locations; short forms are collected here for quick reference.

- **Balance-Sheet Repression** (Section 3.1.2): financial repression implemented through collateral rules, capital treatment, stablecoin reserve rules, custody mandates, and institutional balance-sheet incentives.
- **Administrative Repression** (Section 5.2): the conversion of formally optional financial, identity, compute, and settlement rails into practically mandatory rails through custody defaults, compliance rules, app-store control, tax treatment, institutional mandates, benefit systems, and platform terms of service.
- **Participation Line** (Section 3.6): the household or organizational threshold below which a person or firm lacks the redundancy to act freely across time—to fail, retry, transact, move, learn, refuse coercive terms, or survive shocks.
- **Agency-Preserving Infrastructure** (Section 3.6, Item 9): infrastructure that expands a user's capacity to act without converting the user into a dossier, dependency object, or platform account.
- **Homestead Ratio** (Section 5.6): the share of verified compute, proof generation, and AI-service capacity supplied by open-admission, non-hyperscaler, geographically diverse, independently verifiable operators.
- **Facility Capacity Receipt (FCR)** (Section 17.7): a signed, auditable receipt that extends Facility Energy Receipts with grid, cooling, redundancy, hardware, jurisdiction, and infrastructure-resilience claims.
- **Physical VerifyPrice** (Section 22.8, Appendix A.8): the time, cost, and confidence required to verify the physical infrastructure claims behind a unit of verified work.

- **Wrapper Dominance Ratio (WDR)** (Section 13.6): a reported *pair* — custodial or synthetic exposure share (stock) and wrapper-to-native activity ratio (flow) — never divided into one scalar. A rising WDR indicates the asset may be financializing faster than it is becoming money.

Quick Reference: Physical Sovereignty & Substrate

These terms make the physical layer falsifiable and name the closed-stack competitor. Full treatments appear at the cited locations.

- **Energy & Physical Interdiction** (Section 5.5): adversary class using curtailment, rationing, tariff discrimination, interconnection denial, and load prioritization to raise verification cost without prohibiting cryptography.
- **Sovereign Optionality ($\mathcal{O}_s$)** (Section 17.8): capacity-weighted index from Facility Capacity Receipt fields; feeds risk haircuts and Red Line 13.
- **Disruption-Adjusted VerifyPrice** (Section 17.10.2): probability-weighted verification cost across physical disruption states; resilience is subordinate to constitutional VerifyPrice SLOs.
- **Red Line 13: Energy Sovereignty Failure** (Section 30.6): verification affordability becomes a sovereign policy variable; monitors upstream of Red Line 1.
- **Closed vs. open sovereign stack** (Chapter 33): convergent trust-minimization applied to matter; inverted locus of sovereignty (state vs. participant).
- **Real-Resource VerifyPrice SLO vs. Physical VerifyPrice** (Appendix A.8, glossary): the constitutional cost band on verifying receipts in real resources on reference hardware (Red Line 1) vs. the cost to audit FCR infrastructure claims (Red Line 10).

Quick Reference: Market Realization & Price Formation

This is the market-structure vocabulary for reasoning about price without confusing it with adoption. Full treatments appear at the cited locations.

- **Native instrument vs. external financial wrapper** (Section 8.4): native instruments include the base asset, Work Credits, evidence artifacts, project credit, staking, LP, and corridor claims. Touching the protocol does not make an instrument monetary. External financial wrappers (spot ETFs, ETPs, treasury companies, custodial balances, futures, options, swaps, leveraged ETPs) do not touch it. Unqualified "wrapper" means the latter.
- **Value capture vs. price capture** (Section 13.2.1, Corollary 13.1): value capture routes native demand through the monetary object; price capture is demand for exposure to the asset's price and can occur with no protocol use. Neither implies the other, in

either direction.

- **Market Realization Plane** (Section 13.7): the external financial machinery through which claims on the native asset are represented, financed, allocated, and priced. Orthogonal to the seven-layer stack—a plane around it, never a Layer 7.
- **VerifyFlow** (Section 26.8, Appendix A.9): the fourth verification family, measuring the external financial representation and price-transmission state of the native monetary object.
- **Flow elasticity (ε)** (Appendix H.1): the responsiveness of a wrapper's shares outstanding to changes in its value per share; the observable signature of holder operating rules. Measured, not structural.
- **Net mechanical gain (κ)** (Appendix H.3): $L(L-1)+\varepsilon L^2$, the coefficient converting an underlying return into a wrapper's return-coupled exposure demand. Defined for every L; the rebalancing term $L(L-1)$ is positive for long ($L>1$) and inverse ($L<0$) daily-reset products alike, and vanishes at $L=1$, where $\kappa=\varepsilon$.
- **Wrapper Recycling Ratio (WRR)** (Appendix H.5): $-\varepsilon L/(L-1)$; equal to 1 when holder flows exactly offset gross rebalancing, below 1 when the wrapper amplifies, above 1 when it is countercyclical. Domain $L \notin \{0,1\}$, inverse products included: for unlevered wrappers ($L=1$) the denominator vanishes, the wrapper has no rebalancing trade to recycle, and the published object is ε itself with the label "unlevered" (Appendix H.5).
- **Return-coupled vs. return-decoupled flow** (Appendix H.7): return-coupled flow responds to the return and shapes volatility and persistence; return-decoupled flow arrives independent of it and shapes destination and level. The decomposition is exhaustive and exclusive: the ε-component belongs to the coupled side (inside κ), the intercept α to the decoupled side, and the residual is published unabsorbed (Appendix H).
- **Wrapper–Native Growth Gap (WNG)** (Section 26.8): $\Delta \ln E^{\text{wrapper}} - \Delta \ln U^{\text{native}}$, computed on unit-denominated quantities over a trailing four-quarter window (revaluation stripped); persistently positive means financial exposure is outgrowing monetary use.
- **Compositional adversary** (Section 5.2): a failure mode with no malicious actor, in which locally rational operating rules combine into concentration, procyclicality, recursive leverage, and uninformative prices.
- **Capital Survival Ratio (CSR)** (Section 26.8.4): whether successive investor cohorts in a wrapper preserved capital, as distinct from whether the product preserved AUM.
- **Policy Concentration Ratio** (Section 26.8.3): the share of net asset demand governed by the largest common rebalancing templates, treasury algorithms, or agent policies.

- **Delegation hop, trust perimeter, Cross-Perimeter Share** (Section 7.2): a task specification crossing a principal–agent boundary; the set of processes a party can sanction outside the protocol; and the fraction of hops that cross such a perimeter, which is the demand variable for portable receipts and is not yet computable.
- **Standing and standing inversion** (Section 7.6, Section 5.2): the locus of ownership, consent, liability, and revocation over an acting process; and the failure in which agents satisfy every measured holder condition while human title over them stops binding.
- **Principal Map, HRS, NPCR** (Section 26.8.9): an overlay classifying native demand and holders by ultimate principal; the human-revocable share of agent-controlled native balances; and policy concentration computed over native positions.
- **Automatic vs. marginal buyer** (Section 3.1.3): formula demand versus residual price-setting demand; quantity of bids is not quality of risk absorption.
- **Duration of the claim vs. duration of the project** (Section 3.3): a repression-resistant store of value must not be a coupon; civilization still needs a duration warehouse. Not a layer.
- **Buyer-quality checklist** (Section 26.8.7, Appendix H.14): persistence, horizon, liability match, loss-bearing capacity, countercyclicality—fields, not a scalar.

Part I
Context & Claims

This Part sets the stage: it explains why soft guarantees (central banks, broadcast media, credentialed authority) are failing under debt, AI, and platform control, and why new monetary primitives are needed. It introduces the triad (Privacy, Proofs, Compute), the threat model, the layered "cypherpunk stack" that the rest of the thesis will fill in, and the delegating machine that is increasingly the buyer of all three.

Chapter 1

Introduction

Start Here: Four Reading Paths

This chapter opens with the cultural and historical argument for why anchors are being sought at all. That framing matters, but it is not the load-bearing part of the thesis, and readers who already accept the premise should skip ahead rather than wade through it.

If you allocate capital or underwrite risk: read the *Executive Memo* and *Reader's Map* in the front matter, then Section 1.2 (the eleven premises), Section 8.1 (the argument in one page, with its weakest links named), Lemma 13.1 (the condition that decides whether any of this captures service value), Section 30.6 (the eighteen ways to prove the thesis wrong), and Chapter 16 (the evaluation questions to put to any candidate instrument).

If you build infrastructure: start at Chapter 6 for the seven layers, then Section 24.1 for the twelve primitives and Chapter 24 for the four reference applications. Chapter 31 is the 12–36 month build sketch.

If you are here to find the weak points: Section 5.6 states the boundary conditions no engineering removes, Chapter 34 answers the strongest known objections—including the completeness objection in Section 34.5.1—Chapter 33 concedes what a competent closed sovereign stack does better, and Table 8.1 rates the argument link by link.

If you want the whole claim in one sentence: *Thesis in Plain Language*, in the front matter.

One warning that applies to every path: nothing in this document treats price as evidence. The monetary claim is tested on protocol telemetry, and price is interpreted separately (Corollary 13.1, Section 26.8).

Every monetary epoch begins with an argument about what is real. James Dale Davidson and William Rees-Mogg, in *The Sovereign Individual* [Davidson & Rees-Mogg 1997], argued that the decentralization of computer networks would erode centralized power, and that governments, in their death throes, would reach for more aggressive tools of control: *capital controls, nationalization, outright authoritarianism.*

The old guarantees (central banks, broadcast media, credentialed authority) no longer hold their shape under the pressure of digital networks and foundation models. They are under

attack, not by a single conspirator, but by the mechanics of decentralization itself, which dissolves the monopolies that once defined our shared reality.

Marshall McLuhan [McLuhan 1964] saw an even deeper shift: the move from print, a one-to-many medium that centralizes narrative, to electronic networks, a many-to-many mesh that fragments it. A print-created reality is a centralized reality: whoever controls the press controls the narrative. This underwrote 20th-century monetary supremacy and global settlement currency backed by a monopoly on violence: the post-1971 US petrodollar.

In the 20th century, citizens got their singular cultural feed from *The New York Times* over breakfast, and later their dualistic, pseudo-antagonistic programming from CNN and Fox News. In the 21st, the internet (many-to-many) erodes this centralization of narrative. Take X, once Twitter: authoritarian regimes that attempt to censor information find it very hard to stop leakage. Within minutes, people with phones can produce a preponderance of evidence that either supports or invalidates the prevailing story.

But a countervailing force has arrived: AI. Generated content has crossed the threshold past which an ordinary viewer can no longer tell it from a recording. Synthetic faces, voices, and scenes now compete with direct sensory experience.

This creates a world wealthy in symbols but poor in anchors: a sea of abstractions with no obvious way back to sensible shores.

> So we look for new anchors.

Cypherpunk cryptography is the unifying answer: privacy by default, proof by construction, and compute that only gets paid when anyone can verify it. It replaces authority with protocols and gossip with receipts. In this frame, money is not a promise from a platform; it is what remains when verification is cheap and permission is irrelevant. Cypherpunks do not ask for integrity; they **instrument** it. They do not trust platforms; they price receipts. One precision is worth making immediately, because the slogan this credo is usually compressed into—*don't trust, verify*—inverts what the machinery is for. The object of cheap verification is not perpetual suspicion, which no economy has ever been able to afford; it is the restoration of *safe delegation*. Verification exists so that trust can be extended to what has earned it and withheld from what has not. A system in which everyone must check everything has not solved the trust problem, it has made the trust problem everybody's second job, and it would be a worse arrangement than the one it replaced. The success condition runs the other way: checking becomes so cheap, and so credible a threat, that almost nobody ever has to perform it—which is exactly the no-questions-asked property Section 13.3.10 identifies as moneyness. **The instruction this thesis actually follows is verify in order to trust.** The triad is that credo made economic: privacy that preserves agency, proofs that travel, and compute that earns only when anyone can verify it.

Bitcoin's SHA-256 puzzle [Nakamoto 2008] mints digital scarcity by tying consensus to

thermodynamic cost, and it is worth being exact about why that works: the work has *no external buyer*. Nobody outside the system wants a partial hash preimage, so nobody can subsidize it, mandate it, price it, or withdraw demand for it. The cost is therefore non-negotiable, and the money inherits that objectivity. The design direction explored in this thesis keeps what makes PoW legitimate (open admission and unpredictable leader election) but swaps the work function, so that the "lottery ticket" is earned by producing succinct, publicly verifiable receipts of compute somebody wanted. Miners win by attaching a proof that a market-demanded computation was done correctly.

A natural anchor workload is matrix multiplication. MatMul-PoUW constructions [Komargodski & Weinstein 2025] make verification asymptotically cheaper than naive production (for example, $O(n^2)$ verification vs. $O(n^3)$ multiplication) while the miner's overhead for producing the certificate is $(1 + o(1))$ relative to naive multiplication. That turns AI's core primitive into **verified FLOPs**: a commodity unit anyone can check cheaply. Projects like Nockchain [Nockchain 2025] show the zk-PoW design space operating in the wild (fair-launch ethos, scope-minimized "dumbnet," explicit issuance), illustrating that proof-carrying work can be wired into consensus without appointing gatekeepers. Keep the lottery; change the work. Make the prize a receipt the public can verify—and accept, in exchange, that the receipt has a buyer and the buyer has a jurisdiction.

Caveat: MatMul-PoUW constructions and useful-work consensus are promising research directions, not yet proven production monetary bases. The arXiv literature [Komargodski & Weinstein 2025] describes the optimal-security claim as a *conjecture*; deployment-grade engineering remains ahead. This thesis frames PoUW as the destination of a research program, not a shipping product.

> Operationally, the goal is simple: bind useful work to the hash race **without introducing trust**.

Two deployable patterns are in scope:

- **(A) Hash-gated useful work.** A SHA-family threshold confers short-lived eligibility; a block is valid only if it carries a PoUW artifact (e.g., MatMul proof or zk-proof) seeded by header randomness.
- **(B) Proof-first selection.** Miners race to post useful-work receipts to a mempool; header entropy resolves ties and timing.

Both patterns must prevent precomputation (epoch seeds, commit-reveal), avoid closed-hardware dependencies, and expose verifier-light clients plus dispute/slash routes for junk artifacts.

When blocks routinely carry proofs that clear public SLOs and decentralization telemetry remains healthy (time-to-first-proof, top-N share, geo/ASN spread), block rewards bind issuance to capacity the world already buys: ZK proving, matrix multiplication, verified in-

ference. So **Privacy, Proofs, and Compute may begin to earn monetary premium**—if the conditions outlined in this thesis hold.

State the cost of this at the outset, because it is real and it runs the other way. An external buyer of the work is a party who can be regulated, subsidized, or coerced, and whose demand can therefore be moved by something other than the resource cost of the work. That is precisely the exposure SHA-256 avoids by having no buyer at all, and it is why Bitcoin's work function is the monetarily stronger construction. Proof-of-useful-work is not an improvement on it; it is a different trade—**monetary objectivity exchanged for capacity relevance**—made in the belief that the second is what a dense digital civilization will be short of. The trade is argued at full strength, against itself, in Section 34.1.1, and a reader who wants the concession before the argument should go there first.

The rest of this Part translates that intuition into claims, contributions, a threat model, and a layered architecture.

What Would Falsify This Thesis?

The store-of-value claim fails if any of the following persist without remediation: (1) verification becomes expensive or gated (VerifyPrice breaks); (2) refund safety fails on admissible corridors; (3) verification concentrates in a single hardware profile or jurisdiction; (4) telemetry is captured or becomes unverifiable; (5) fee coverage collapses and workload demand becomes purely speculative; (6) users can consume triad capacity without touching the native asset; (7) wrapper and custodial exposure grows while native usage stagnates; (8) physical infrastructure claims become unverifiable; (9) AI compute enclosure persists; (10) lawful privacy collapses socially; (11) verification affordability becomes a sovereign policy variable; (12) the capturable wedge falls below the protocol take; (13) a native-collateral drawdown drives operator exit, declining Delivered Verified Capacity, released float, and unrecovered service SLOs; (14) no regime-contingent holder-side flow is ever measurable on the base asset, or a balance held in advance does no better than Bitcoin converted on the day; (15) cheap verification fails to narrow how far independently underwriting lenders disagree on haircuts against base-asset positions, relative to opaque comparables of the same volatility; or (16) protocol insiders themselves decline to denominate their short-tenor obligations in the unit. These are illustrative and compress several red lines into one clause; the full set of eighteen red lines is developed in Chapter 30.

1.1 Central Claims

The thesis argues five things.

1. **Post-Bretton Woods money relies increasingly on compliance infrastructure.**

As convertibility receded, taxation capacity, legal tender status, central-bank balance sheets, collateral rules, and the compliance perimeter became the practical support structure of fiat money. When compliance becomes weaponized and asset freezing becomes policy, neutral stores of value become necessary infrastructure, not ideological luxuries.

2. **Three cryptographic capacities may support a monetary candidate under measurable conditions.**
   Privacy (censorship-resistant settlement that preserves agency), **Proofs** (portable attestations of computation and provenance), and **Compute** (useful work wrapped in succinct guarantees) are distinct services. They support a monetary candidate only if their service path remains stress-deliverable, their demand cannot bypass the candidate asset, and native holders warehouse its residual risk. They need not share one token.
3. **A modular stack can operationalize this triad.**
   Four reference applications (private treasury & payroll, media provenance, verified inference, proof/compute procurement) are built on twelve primitives, with **verification asymmetry** and **VerifyPrice** as the key economic metrics that turn proofs and verified FLOPs into commodities rather than platform IOUs.
4. **Verifiable machines (Layer 0) are essential for physical-world claims.**
   For pure proof correctness, sound cryptography should tolerate untrusted provers. Layer 0 becomes essential when the claim includes facts about physical capture, energy use, machine identity, witness confidentiality, side-channel resistance, or useful-work fairness. Open or sampled hardware is not ornament; it is the base layer that makes physical-world claims credible.
5. **Accrual is not moneyness, and price is not adoption.**
   Fees, burns, and collateral route service value to a base asset, and what they establish is a cash-flow claim and a collateral lockup. Moneyness requires mechanisms a cash-flow valuation does not reach. Price, meanwhile, can be set by wrappers, leverage, and allocation rules with no native use at all. Both separations are instrumented rather than assumed, and each is where a weaker version of this argument would have hidden its conclusion. A third follows from them: holding the asset is worth something the service is not only if a held balance beats Bitcoin converted on the day, and that is tested, not assumed (Section 13.1).

The thesis begins from a plural answer: the reserve question is not answered by one object, and the assets that already answer parts of it are not superseded by this one. Alongside them, a dense digital civilization needs a **triad of verifiable necessities**: Privacy, Proofs, and Compute. Treat this triad as services that may support an associated cypherpunk base asset:

- **Privacy**: the right to hold and move value without chokepoints (to preserve agency).
- **Proofs**: cryptographic attestations of origin, integrity, identity, or computation (to crystallize truth).
- **Compute**: useful work (ZK proving, matrix multiplication, verified inference) that applica-

tions demand and that blockchains can verify cheaply (to power intelligence).

These are not slogans. They are the cypherpunk stack distilled into services: what cannot be forged, what no one has to bless, and what everyone can check. They are *scarce capacities* the world may continually buy because life in a dense, digital civilization requires them. Each has its own economy. Whether one monetary asset should sit beneath all three is a hypothesis tested below, not a conclusion smuggled into the definition.

Institutionally, this is not just a design for another chain. It is a **research and engineering agenda**: a Bell Labs for proof-of-useful-work and the zk economy. Where the original Bell Labs turned Shannon's information theory [Shannon 1948] into cables, switches, and semiconductors, the mandate here is to *turn privacy primitives, zero-knowledge, and verifiable compute into everyday infrastructure*: receipts, rails, and verifiable machines that other people build on without thinking about it. This document is the charter for that lab.

Each primitive is designed to be credibly scarce, permissionless, censorship-resistant, and cheap to verify. Gold condenses geology; Bitcoin condenses thermodynamics; this triad condenses the utilities of the information era into assets: things that cannot be faked and do not ask permission, cheap for anyone to verify yet costly to produce or deny. Three different hard facts, answering three different questions, none of them a later edition of another (Section 34.1.1).

> **The claim under test.** Under sustained administrative and financial repression, a bearer asset may earn monetary premium if holding it preserves private settlement, portable proof, and access to verified compute after ordinary substitutes weaken. The full service path must remain stress-deliverable; demand must not bypass the asset; a persistent, self-custodied, loss-bearing holder constituency must absorb residual financial risk; and project credit must remain separate. Whether any network satisfies those conditions—or whether the triad should share one asset—is an empirical question. This thesis does not ask to be believed. It asks to be measured, and it names in advance every reading that would retire it.

The thesis will test this through a conditional chain: utility demand creates standardized work; standardized work produces receipts; receipts represent stress-deliverable service; those receipts enable markets; markets produce fee flows; fee flows support a scarce asset *only under non-bypassable accrual rules*; a persistent, self-custodied, loss-bearing, regime-responsive holder constituency anchors that asset; and only then, if the asset remains liquid, neutral, verifiable, legally holdable, and agency-preserving, may it earn a store-of-value premium.

1.2 Premises

The argument below is the steel-man form used on the project site and in outreach. Each premise is developed in the body; together they bound what the thesis does and does not claim.

P1. Soft guarantees are weakening. Modern money, media, and infrastructure increasingly depend on compliance systems, custodians, platforms, hardware vendors, and institutions whose guarantees are politically and technically fragile (Chapter 2, Chapter 3).

P2. The digital economy has three unavoidable needs. A dense AI civilization needs private settlement, portable attestations, and verified compute—Privacy, Proofs, and Compute as scarce capacities, not slogans (Section 1.1, Part II).

P3. These needs can become verifiable commodities. If workloads are standardized, outputs become receipts, and verification stays much cheaper than production, proofs and verified compute can trade as public facts rather than platform promises (Chapter 22, Appendix A).

P4. Gross capacity is not deliverable service. Installed power and hardware are inputs, not usable output. Monetary-grade capacity is the scenario-specific surviving flow through energy, conversion, cooling, hardware, network, workload, proof, verification, settlement, and user access (Section 8.2, Section 17.9).

P5. A store of value requires more than utility. Indispensable services do not automatically become money; demand must accrue to a scarce asset rather than leaking to operators, hyperscalers, custodians, or fiat rails (Chapter 4, Section 34.1.5).

P6. Value capture requires enforceable monetary design. The asset must be required for core fees, partially burned or retired, staked as collateral, issued under discipline, and protected from bypass channels (Value Capture Lemma, Section 30.6).

P7. Gross native demand is not a monetary anchor. Just-in-time fee acquisition, burns, operator inventory, collateral requirements, and wrappers do not by themselves create a persistent, self-custodied, loss-bearing constituency that adds through stress (Section 26.8.8, Section 26.8.7).

P8. The system must remain falsifiable. Verification cost, reachability, settlement safety, decentralization, fee coverage, native-asset capture, delivered capacity, holder quality, and agency must be public telemetry—not marketing claims (Chapter 26, Section 30.6).

P9. Market price is not proof of monetary adoption. Even when the protocol loop works, observed price can be set by wrappers, leverage, dealers, and allocation rules without native use (Corollary 13.1, Section 26.8). VerifyFlow instruments the second loop; it does not replace VerifyPrice, VerifyReach, VerifySettle, or the Native Monetary Buyer Map.

P10. Duration-neutral money is not duration finance. A repression-resistant store of value cannot be a coupon the state can pin. Civilization still needs someone to warehouse the years between pouring concrete and producing power. Project notes are explicit loss-bearing

credit, not monetary backing (Section 3.3, Section 34.5.2, Section 34.5.1).

P11. **The buyer of verification is increasingly a delegation.** Machine processes acting for principals several hops away cannot do due diligence at speed, so verification is the only reliability mechanism that fits inside the hop. That demand is structural only where delegation crosses trust perimeters and the work is too consequential to accept unchecked; its coercibility depends on who operates the delegating agent; and whether a human still holds standing over that agent is a question the stack must answer rather than assume (Chapter 7, Section 5.2).

> **Conclusion (conditional).** A base asset is only a monetary candidate. It may earn premium only if the seven-layer stack delivers usable service under stress, demand cannot bypass it, and a persistent self-custodied constituency bears loss through the relevant regime. Work Credits remain typed capacity or service claims; FCR, FER, and PIDL artifacts remain evidence; project notes remain duration-bearing credit; and LP or staking positions remain derivatives or operating claims. No useful service is promoted to money by terminology.

1.3 Contributions

The thesis makes eight contributions. The first two are the engineering substrate; the next four are where the argument is new; the last two are what makes it testable and buildable.

1. **A threat model that assumes repression, and a seven-layer stack built against it.** The adversary is intentional: financial and administrative repression, censorship and shutdowns, hardware and identity capture, platform-mediated reality, and energy interdiction that raises the cost of verification without prohibiting cryptography. Two failure kinds have no adversary at all: the compositional adversary, in which locally rational operating rules combine into concentration and uninformative prices (Section 5.2), and standing inversion, in which agents satisfy every measured condition while human title over them stops binding (Section 5.2). Against both, the thesis specifies seven layers, from matter to institutions:
   - Layer 0: Verifiable Machines & Energy
   - Layer 1: Reachability (communications & transport)
   - Layer 2: Distribution & Execution (software supply & runtime)
   - Layer 3: Identity & Claims (humans and machines without doxxing)
   - Layer 4: Truth & Work (proof systems, PoUW, VerifyPrice)
   - Layer 5: Value & Settlement (privacy rails, non-custodial flow)
   - Layer 6: Governance & Telemetry (keeping neutrality and resilience measurable)
2. **Verification asymmetry, priced.**

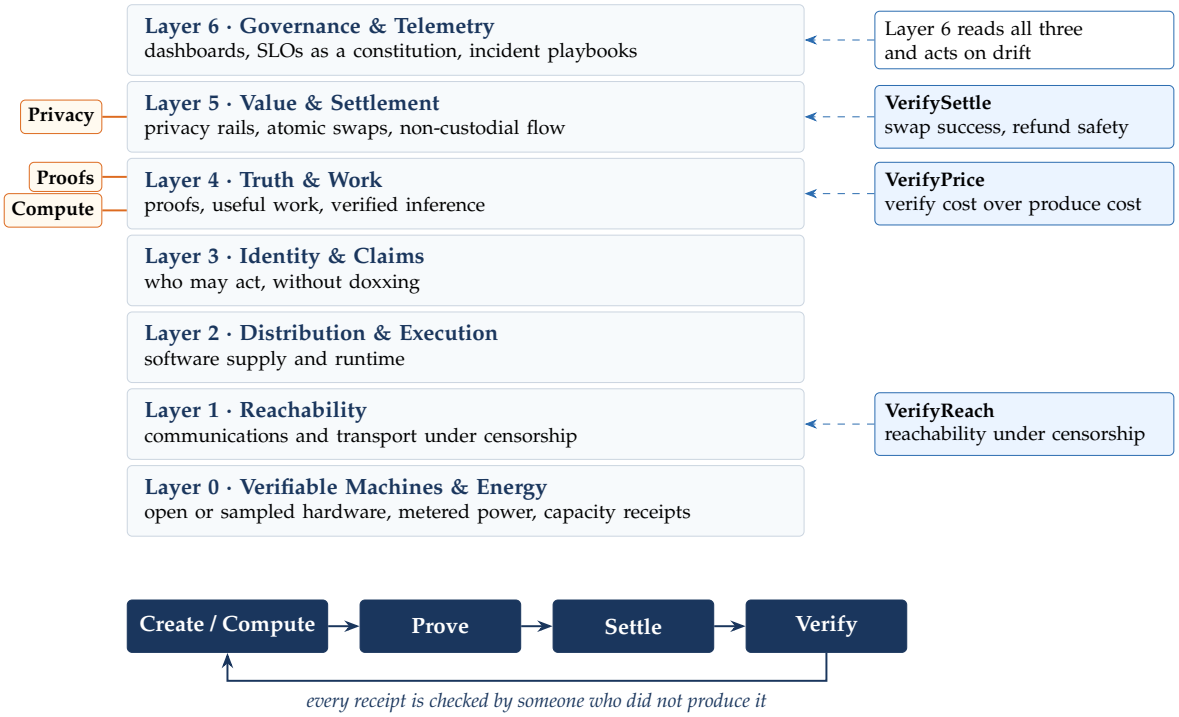


Figure 1.1: Architecture map: the seven-layer stack, the Create/Compute → Prove → Settle → Verify loop, the triad on the layers that supply it, and the three telemetry families.

The ratio

$$r(W) = \frac{v(W)}{p(W)}$$

between verification and production cost for a workload W is formalized, and **VerifyPrice**, a public vector of p50/p95 verification times, costs, and failure rates on reference hardware a person can own, is made the hinge on which proofs and verified FLOPs become commodities rather than platform IOUs. It is denominated in real resources, so it is exogenous to token price and not to the conditions under which power and hardware can be obtained.

3. **Deliverability as a flow, not a stock.**
   The **Topological Scarcity Lemma** replaces nameplate capacity with the stress-deliverable flow through the narrowest non-substitutable edge between energy and holder utility, and **Delivered Verified Capacity** measures it by scenario (Section 8.2, Section 17.9). Issuance is bounded by that flow. The associated **Pressure–Capacity Corridor** predicts an inverted U: little differential value while institutions work, the strongest case while substitutes weaken and the open stack still runs, and a falling one when pressure disables the stack or the holder.
4. **Accrual separated from moneyness.**
   The reference design's value-capture channels (Value Capture Lemma) establish when service demand accrues to its base asset, and the fee-incidence argument concedes what accrual buys: a competitively priced cash-flow claim and a collateral lockup, both reproducible by a discounted-cash-flow valuation (Section 13.3). The monetary claim is then carried by three mechanisms such a valuation does not reach: **pledgeability without due diligence**, reached by symmetric knowledge rather than the symmetric ignorance every prior money-like asset relied on; **denomination**; and a state-contingent **holder-side service flow** (Section 13.3.10, Section 13.7.3, Section 13.3.11). Each has an instrument and a red line. None is sized. The third is tested first against its strongest alternative, Bitcoin converted on the day: what holding buys is a balance that stays spendable when acquisition fails, and an access experiment decides whether that is ever worth its risk (Section 13.1).
5. **Price separated from adoption.**
   The **Market Realization Plane** names the external machinery—wrappers, custody, leverage, dealers, and allocation rules—through which a native asset is represented and priced, and **VerifyFlow** instruments it: flow elasticity, net mechanical gain, the Wrapper Recycling Ratio, and the Wrapper–Native Growth Gap (Section 13.7, Section 26.8). The same discipline reads the reference asset of the incumbent system, whose marginal demand is increasingly commissioned by benchmark rules and delivered by leverage (Section 3.1.3). Price is something to explain, not evidence to cite.
6. **A holder anchor, and a model of who is buying.**
   Gross native demand is not a monetary anchor. The **Native Monetary Buyer Map** asks whether a persistent, self-custodied, loss-bearing constituency holds through stress (Sec-

tion 26.8.8). The **delegation economy** models the marginal buyer of verification as a machine delegating across trust perimeters (Chapter 7), and the **Principal Map** classifies holders by ultimate principal, so that a cohort which is one operating policy at many addresses, or which no human can still revoke, is not counted as an anchor (Section 26.8.9).

7. **A modular stack: one loop, twelve primitives, four reference applications.**
   The **Create/Compute Prove Settle Verify** loop is instantiated as private treasury & payroll, media provenance & authenticity, verified inference, and proof/compute procurement, built on a reusable kit: a Proofs-as-a-Library SDK (PAL), a Privacy Rails Kit (PRK) for non-custodial settlement, the PIDL receipt schema, MatMul-PoUW and verified-inference harnesses, canonical workload registries, multi-ZK adapters, SLA escrow and slashing, neutral routers, bridge-safety templates, and the telemetry that keeps useful work and neutrality measurable across chains and vendors.
8. **Falsification as architecture, and an institution to run it.**
   Eighteen red lines, each named, measurable, and able by itself to retire or downgrade the thesis at the severity its box states, are read on public boards with SLOs as the constitution (Section 30.6, Section 26.9). **VerifyReach** and **VerifySettle** extend VerifyPrice to reachability and refund-safe settlement. An adoption model, an operator and investor checklist, and implementation sketches make the stack buildable over the next 12–36 months, and **Kardashev Labs** is specified as the build-and-measure institution, with red lines of its own (Chapter 32).

1.4 End-to-End Vignette: The Loop in Action

Before the theory, a concrete story that walks the **Create/Compute Prove Settle Verify** loop. This vignette shows how the stack works as a felt reality, not just a framework.

Scenario: A Company Runs Private, Auditable Payroll

TechCo is a software company with employees and vendors across several jurisdictions. Its CFO needs to run payroll while preserving salary confidentiality, avoiding unnecessary exposure of the transaction graph, and giving auditors and tax authorities scoped evidence that payments were authorized, complete, and properly reported.

Step 1: Create (Intent) The CFO opens TechCo's treasury dashboard (built on PRK, the Privacy Rails Kit). She creates a payroll batch:

- 47 employees
- Total: 142,000 USD-equivalent, quoted in dollars and settled in the base asset, with corridor recipients paid out in XMR

- Policy: "Salaries are private; aggregate spend is auditable; individual amounts disclosed only with employee consent."

The dashboard compiles this into a **claim**: "Pay these 47 recipients the specified amounts, under this policy, by end-of-day Friday."

Step 2: Prove (Compliance + Privacy) The claim is sent to PAL (Proofs-as-a-Library). PAL compiles it into two proofs:

1. **Compliance proof:** A ZK proof that the batch satisfies TechCo's internal policy (aggregate under budget, recipients are on the approved list, no single payment exceeds threshold). This proof reveals *nothing* about individual amounts or identitiesonly that the rules were followed.
2. **Integrity proof:** A hash commitment binding the full payroll data. This hash is stored on-chain; the actual data stays encrypted in TechCo's vault.

Both proofs are packaged into a **PIDL receipt**: claim hash, proof hashes, workload ID, SLA tier, timestamps, and a prover signature.

VerifyPrice checkpoint: The compliance proof took 2.3 seconds and cost $0.004 to generate, paid to the proving factory in the base asset (the required fee medium of Lemma 13.1). Independent verification takes 0.8 seconds on a laptop. This is well within the SLO ($p_{95,t} \leq 5$s, $p_{95,c} \leq \$0.01$).

Step 3: Settle (Private Execution) With proofs in hand, PRK executes the payroll:

- For employees in jurisdictions with limited banking access: atomic swap via the **BTC↔XMR corridor**, paying the corridor fee in the base asset. TechCo's BTC is swapped for XMR, which is sent to employee wallets. The corridor is non-custodial; if anything fails, every party has a bounded-time no-loss exit (refund-safe, Section 23.4).
- For employees elsewhere: direct settlement in the base asset over its **shielded pool**. Each payment is encrypted; only the recipient, and TechCo through per-epoch disclosure, can see the details.

What this vignette shows, and what it does not. It shows the stack closing end to end. It is not evidence for the monetary claim, and its own denomination says why: the batch is *quoted* in dollars and *settled* in the unit, which is the payment-rail configuration that Red Line 18 names as the null and that bitcoin has occupied for seventeen years (Section 13.7.3). The base asset appears here as the required fee medium and the settlement unit — the routing rules of Section 8.4, which establish accrual and nothing more. Whether TechCo ever writes a contract *in* the unit is the question the thesis measures rather than assumes.

VerifySettle checkpoint:

- Swap success rate: 98% (2 retries due to network latency, both succeeded; counted in the 98%, so 49 of 50 corridor payments landed within two attempts).

- Refund safety: 100% (all potential failures would have returned funds).
- Time to finality: median 4 minutes, p95 11 minutes.
- Anonymity set: 12,000+ active notes in the shielded pool.
- These figures are illustrative of the checkpoint's shape, not observed telemetry; the published VerifySettle schema is the same one the Settlement & Privacy Board reports (Section 26.9).

Step 4: Verify (Audit Trail) After settlement, the following are publicly verifiable:

- **On-chain:** The PIDL receipt exists, the proofs verify, the aggregate amount was transferred at the specified time.
- **By TechCo's auditors:** Using viewing keys, auditors can see the full breakdown (who got paid how much) and confirm it matches the compliance proof.
- **By employees:** Each employee can verify their own payment arrived, using their private key.
- **By anyone:** The compliance proof demonstrates policy was followed, without revealing any private data.

What the adversary sees:

- An external observer monitoring Country A sees: "TechCo moved some BTC into an XMR corridor." They cannot see amounts, recipients, or purposes. TechCo can provide auditors and tax authorities scoped viewing keys and PIDL receipts demonstrating that payments were authorized and properly reported.
- A competitor monitoring the blockchain sees: "Some shielded transactions occurred." They learn nothing about TechCo's payroll structure or employee compensation.
- A hacker who compromises a single employee's device sees: that employee's payment. They cannot reconstruct the full payroll or identify other employees.

What TechCo gains:

- **Privacy:** Payroll data is not exposed to competitors, data brokers, or unrelated intermediaries by default. Auditors and tax authorities receive scoped evidence via viewing keys.
- **Compliance:** Auditors get cryptographic proof that policy was followed, without needing to trust TechCo's word.
- **Resilience:** Even if banks freeze accounts or exchanges delist, the privacy corridor remains operational.
- **Receipts:** Every step produced a verifiable artifact (PIDL receipts, proofs) that can be archived, audited, or used in disputes.

Stage	What Happens	Output
Create	CFO defines intent + policy	Claim
Prove	PAL compiles compliance + integrity proofs	PIDL receipt
Settle	PRK executes via privacy corridors	Value transferred
Verify	Anyone can check proofs; auditors use viewing keys	Audit trail

Table 1.1: The Create/Compute Prove Settle Verify loop in the payroll scenario.

The loop in summary: This is **lawful privacy**: default-private, optional-disclosure, with receipts that anyone can verify. The triad (Privacy for settlement, Proofs for compliance, Compute for proof generation) works together to make the flow possible.

Variation: Media Provenance

The same loop applies to a journalist publishing a video:

1. **Create:** Camera with secure element captures footage, emitting a signed provenance attestation (device ID, timestamp, geolocation hash).
2. **Prove:** Edit suite issues proofs of each transformation (crop, color grade, caption). Each edit is a new PIDL receipt referencing the parent.
3. **Settle:** (Optional) If the video is monetized, payment flows over privacy rails to the creator, with receipts linking payment to provenance.
4. **Verify:** Any viewer or platform can verify the chain: "This video originated from camera X at time T, was edited as follows, and has not been tampered with since." Deepfake detectors can check against registered provenances.

Platforms that strip metadata cannot strip the on-chain PIDL receipt, but persistence is not discoverability: a re-encoded copy no longer hashes to the receipt, and the stripped copy is matched back to it by the soft binding of Section 20.7 — a perceptual hash registered at capture — with the analog hole as the stated residual. The proof persists even if the platform delists the content; whether a given copy can be led back to it is the registry's job, and its limits are stated where the mechanism is.

These vignettes are not forecasts. They are the concrete scenarios the rest of this thesis is engineered to support. The stack exists to make these flows **cheap, verifiable, and non-custodial** under adversarial conditions.

Chapter 2

The Failure of Soft Guarantees

The twentieth century ran on soft guarantees: the promise that money would preserve purchasing power; that televised images would correspond to facts; that institutions would be slow, sober, and legible. The twenty-first runs on different physics. Monetary systems must service accumulated promises faster than real productivity can grow. Media systems can synthesize plausible realities faster than editors can adjudicate them. Intelligence, once gated by human attention, now scales with silicon. In that environment, guarantees that depended on slow institutions and centralized editors begin to fray.

One soft guarantee is failing without fraud and without default, which is why it belongs in this list rather than in a chapter about mismanagement: the behavioral promise that discipline converts into advancement — gain skills, borrow for education, save regularly, buy the index, purchase a home, and compounding delivers security. Every contract in that recipe can be honored exactly as written while the mapping fails, because the incumbent's balance-sheet wealth is the entrant's acquisition price, and the financing regime that produced the earlier returns has already closed. A trust failure that survives audit is the hardest kind to repair, and it feeds the participation line directly (Section 3.6).

Since the end of Bretton Woods, *compliance*, not convertibility, has been the hidden backing of money. When compliance becomes weaponized and asset freezing becomes policy, savings need privacy the way ships need hulls.

> *When the **symbols of trust** are cheaper to counterfeit than to maintain, the market seeks a new substrate.*

This is not moral decline; it is entropic arithmetic. When verification is costly and authority is centralized, cheaters arbitrage the gap. When verification becomes cheap and public at scale, the equilibrium flips. Our civilization is now crossing that threshold.

"Money is memory with consequences."

It records who did work and who is owed work, then carries that record across time, space, and politics. Gold solved this with geology; fiat with law, taxation, and war; Bitcoin with thermodynamics and code. Each regime answered the same question (what counts as real work, and how do we know?) with a different ontology.

2.0.1 A Prior Transition: What the Press Actually Changed

Each of those regimes also rested on a medium, and each medium restructured the trust that ran on it before anyone noticed [McLuhan 1962]. This is the pattern the media-transition literature [McLuhan 1964] documents—as mechanism, not measurement—and it is worth one paragraph here because the thesis's claim about attestation has a predecessor with a five-hundred-year record.

Manuscript culture anchored trust in persons: the scribe was known, the chain of copying was personal, and verification meant asking a community that remembered. The printing press collapsed the cost of *replication* and, in doing so, moved the anchor from person to artifact: the imprint, the uniform edition, the publisher's mark. What a reader could check changed—from reputation to internal consistency and corroboration across copies—and the standardization that made every letter identical to every other letter of its kind later gave industry its template for interchangeable parts (Eisenstein (1979)). The broadcast era then re-centralized attestation: one attester, many receivers, and verification deferred back to the institution.

The networked era is the first with no native anchor at all. Anyone can publish; nothing arrives pre-vouched. The incumbent answer is the platform label—attestation by intermediary, and therefore by whoever pressures the intermediary. The answer this thesis tests is attestation by proof: the anchor moves from person (manuscript), to artifact (print), to computation itself. The axis is not new; what is new is that the checking is native to the medium rather than borrowed from the institutions that grew up around the previous one.

One further shift belongs in the sequence. In every prior regime the attester persisted. Manuscript culture could ask the scribe; print could sue the publisher; broadcast could summon the network to a hearing. A software agent that produced a computation an hour ago may not exist now, may have been one of ten thousand instances, and may have no operator willing to answer for it (Chapter 7). The attester is ephemeral for the first time. The consequence for this thesis is that the receipt has to outlive its producer: it must be self-verifying, content-addressed, and anchored somewhere the producer's disappearance cannot reach, which is what the anti-memory-hole requirements of Chapter 26 already specify and what this observation gives them a reason for. The media-transition reading, offered as a reading, is that this is re-tribalization in monetary form: an oral culture, in which speakers do not persist, with the communal memory externalized into proofs because there is no one left to ask.

Two cautions keep this framing honest. The press's consequences arrived over centuries that nobody planned, which is an argument for instrumenting rather than forecasting, not an argument that any particular consequence follows. And the literature itself is synthesis rather than experiment; it is used here for pattern, not for quantities.

Memory alone is not money; the record must attach to a scarce object or claim whose holders capture economic value. The triad thesis must therefore show not just that these capacities produce records, but that a specific asset captures the economics of those records.

Our era answers: **work is what machines can do and humans can verify cheaply**. The new memory must therefore be backed by capacities essential to digital life and secured by proofs anyone can check. The unit must survive not because a state decrees it, but because the network will keep buying its utility through booms, busts, and repression cycles.

> "A store of value is not a museum; it is a power plant that stays online."

The same forces that erode the old legitimacy (AI, ubiquitous networks, programmable markets) also supply the antidote. Cryptography can wrap data, identity, and computation in attestations that travel; privacy tech can re-sacralize the personal and the intimate; verifiable compute can make expensive work legible to cheap verification. The system that dissolves trust also furnishes the reagents to precipitate it again, this time as math, not authority. Verification, not violence, becomes the final arbiter of bounded claims about origin, computation, and settlement.

The rest of this thesis makes that sketch precise. It starts from a threat model that assumes repression, not benevolence: debt stocks that make financial repression arithmetically attractive; a communications stack that can be filtered or shut off; hardware and identity systems that can be integrated into surveillance infrastructure. Against that backdrop it proposes a layered architecture, from verifiable machines and energy up through communications, distribution, identity, truth, settlement, and telemetry, and show how three capacities (Privacy, Proofs, and Compute) can be engineered across those layers as services capable of supporting a conditional base-asset monetary candidate. The task is not to find a new story to believe in, but to build and measure the full chain.

During the transition, old-world hard assets—gold, silver, and other geologically scarce commodities—remain rational hedges against the very soft-guarantee failures described here. They require no counterparty, no network uptime, no software update. The question is not *whether* to hold them, but *when* the new stack becomes robust enough to migrate. Until VerifyPrice dashboards are live, privacy corridors are refund-safe, and receipt volume is nontrivial, the bridge assets remain load-bearing. This thesis is not a call to abandon them prematurely; it is a blueprint for what comes next.

Chapter 3

The World Forces New Monetary Primitives

The forces doing the forcing run on two channels at once, and they are the same channel. In truth, the incumbent order is broadcast: one attester—state, rating agency, platform, exchange—asserts, and N receivers accept. In value, the same shape: central-bank credit is one issuer to N accounts, with the sovereign acting as the typesetter of uniform value, every unit identical because one authority struck it. Proofs and bearer settlement are both $N \leftrightarrow N$: anyone checks anyone's claim; anyone settles with anyone directly. The thesis is one observation wearing two dresses—truth and value move on the same topology, and both are moving from broadcast to network. That is a way of reading what follows, not an additional premise: the chapters that follow treat macro pressure (Section 3.1) and epistemic pressure (Chapter 2) as one phenomenon observed from two sides, and each must still carry its own evidence.

3.1 Macro Playbook: Debt Repression Flight to Neutrality

Total global debt (private + public) according to the IMF's latest Global Debt Database [IMF 2025] now sits at a roughly ∼2.4Œ multiple of world GDP [Gaspar et al. 2025], with some economies (such as Japan) exceeding 400% of GDP when public and private debt are combined. Standard scenarios from multilateral institutions point to global public debt approaching ∼100% of GDP by decade-end [IMF 2026], after already exceeding $100 trillion in the mid-2020s [IMF 2024].

Historically, such overhangs are closed via **financial repression**: inflation plus regulation that imposes negative real returns, including overt capital controls and captive-balance-sheet rules, pushing savers toward neutral, censorship-resistant rails. "Voting with your feet" is harder when the exits are guarded.

The Bondholder's Trap (How Repression Works)

Definition. Financial repression is a stealth tax on savings: cap nominal rates, let inflation run, and force captive balance sheets to hold the paper. The result is systematically

negative real returns for depositors and bondholders; an engineered wealth transfer to the sovereign. In the post-WWII advanced economies, real rates were negative in roughly *half the years from 1945–1980*, and the "liquidation tax" (interest-expense savings on the public debt) averaged on the order of 1–5% of GDP per year [Reinhart & Sbrancia 2015].

How the cap is enforced. Overt or implicit yield-curve control (YCC), administered rate ceilings, capital/liquidity rules that force government paper into banks, insurers, pensions, and collateral boxes, plus capital controls to slow leakage. In the U.S. WWII episode, the Fed capped T-bills at 3/8% and long bonds at 2½% until the 1951 Accord [Rose 2021]; when the cap ended, long-duration holders took capital losses as yields normalized.

Why now. With global debt at several times world GDP, public debt ratios trending toward ∼100% of GDP by decade-end, and interest costs compounding, repression is arithmetically attractive to policymakers versus explicit default or politically costly austerity.

Two equations carry the mechanism:

1. Repression wedge (real carry):

$$r_{\text{real}} \approx \frac{1 + i_{\text{capped}}}{1 + \pi} - 1$$

If $i_{\text{capped}} = 2.5\%$ and inflation $\pi = 6\%$, the real return is

$$\frac{1.025}{1.06} - 1 \approx -3.3\% \text{ per year.}$$

2. Duration burn (when the peg breaks): For a bond with duration D, a yield move Δy produces a price move

$$\Delta P \approx -D \cdot \Delta y.$$

A 10-year with $D \approx 8$ takes about -24% on a +300 bps repricing.

A worked example: three years at -3.3% real/yr is $\approx -9.6\%$ real loss compounded; then a +300 bps jump deals another $\approx -24\%$ price loss, roughly -31% cumulative in purchasing-power terms. Each instrument in it is chosen by someone. What they produce together needs no motive to be real: the bondholder's real wealth moves to the sovereign balance sheet, and the sovereign buys time.

3.1.1 Then and Now

The two equations have been run before, at national scale. From 1942 to 1951 the Federal Reserve pegged bills at 3/8% and capped long bonds at 2½% (Fig. 3.1); across the advanced economies, real short rates were negative in roughly half the years from 1945 to 1980 (Fig. 3.2). The bondholder was not defaulted on. The bondholder was held in place while the real value of the claim was taken, and when the Accord released the cap, the long holder paid the second equation as well.

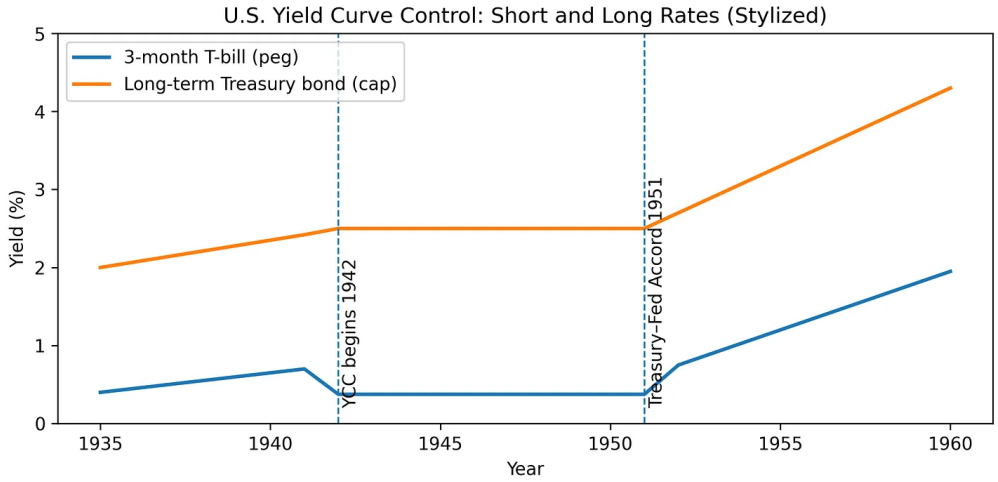


Figure 3.1: U.S. Treasury bill peg and long-bond cap under WWII-era yield-curve control (stylized). Peg at 3/8% on bills, cap at 2½% on long bonds; dashed line marks the 1951 Treasury–Fed Accord. Source: Rose (2021), Federal Reserve Bank of Chicago.

What has changed since is the enforcement. As convertibility receded, the compliance and regulatory perimeter became a larger part of what makes fiat enforceable, especially under capital controls, sanctions, and financial repression. The wartime peg needed a captive banking system. The modern version needs captive balance sheets, custody defaults, and the power to freeze, and the freezing of Russia's reserves in early 2022 showed that the last of these extends to sovereigns. Some investors and institutions have read that episode as a reason to reassess neutral collateral and privacy-preserving settlement. When repression becomes the spread, capital rotates quietly toward neutral, bearer-like rails.

3.1.2 The Treasury Market as Control Panel

Modern repression does not need to announce itself as yield-curve control. It can be routed through balance-sheet architecture. A sovereign that cannot fund itself at market-clearing real yields can manufacture demand for its paper by changing the rules under which banks, insurers, pensions, money-market funds, stablecoin issuers, and custodians hold collateral.

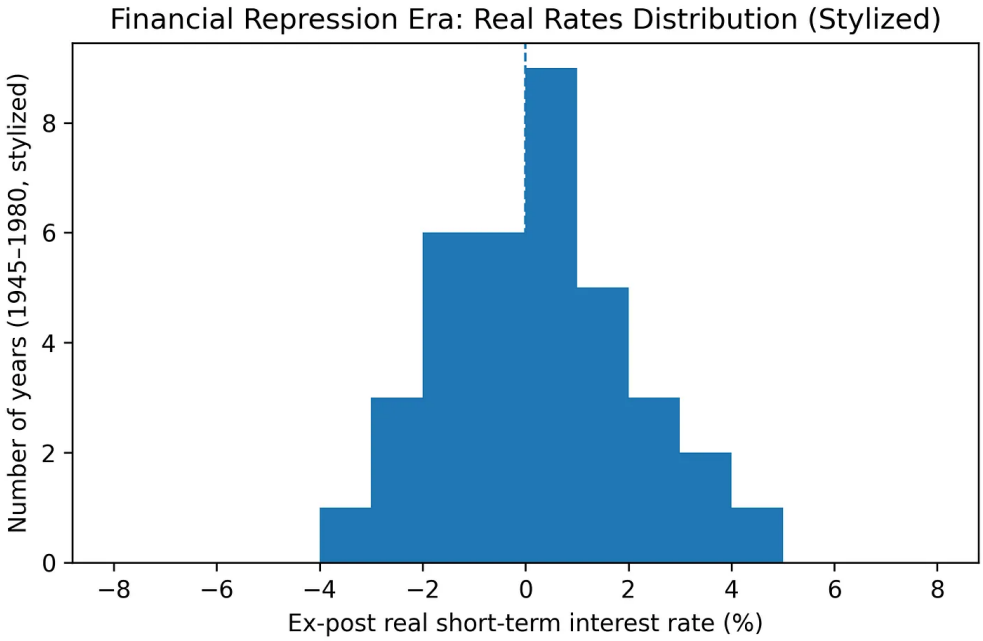


Figure 3.2: Stylized distribution of ex-post real short-term interest rates in advanced economies, 1945–1980. Real rates are negative in roughly half the years, consistent with Reinhart & Sbrancia's estimates of the financial-repression era.

The playbook is subtler than confiscation: lower front-end rates, pressure or tolerate higher long-end yields, steepen the curve, loosen bank leverage treatment for sovereign paper, and let regulated balance sheets absorb debt that private investors would not hold at prevailing real yields. The central bank can claim discipline because its own balance sheet is shrinking, while the banking system performs the functional equivalent of QE. In practice this shows up as redefined high-quality-liquid-asset treatment, eased leverage-ratio exemptions for Treasuries, incentives for banks to intermediate sovereign debt, privileged government collateral in payment and clearing systems, steered stablecoin reserve composition, and regulated custody as the default gateway for institutional capital.

Balance-Sheet Repression

Financial repression implemented through collateral rules, capital treatment, stablecoin reserve requirements, custody mandates, and institutional balance-sheet incentives, rather than through overt yield-curve control or capital controls.

For a post-fiat monetary stack, this matters because the adversary is not only censorship. It is **substitution**. Regulated wrappers, approved custody, stablecoin reserve mandates, and institutional balance-sheet rules can direct capital into price exposure while starving protocol-native privacy, proof, settlement, fee-burn, and collateral loops. The stack must therefore mon-

itor not only whether users *can* access the asset, but whether they are using the monetary rail or merely holding a compliant shadow claim. Section 13.6 formalizes this substitution risk as the **Wrapper Dominance Ratio**.

3.1.3 Who Actually Buys the Paper

A control panel is only as useful as the hands available to absorb what it issues. The composition of those hands has changed, and the change matters more than the level of any particular yield.

Foreign official reserve managers—the patient, price-insensitive creditors of the previous era—have not been dependable net buyers of U.S. government debt for some years. A growing share of marginal demand now comes from leveraged relative-value funds that hold cash Treasuries against short futures positions and finance the difference in the repo market. Federal Reserve staff research [Barth et al. 2025] estimates that Treasury holdings of Cayman-domiciled hedge funds rose by approximately $1 trillion after 2022, reaching roughly $1.85 trillion by the end of 2024, financed through bilateral and centrally cleared sponsored repo.

Two qualifications are load-bearing, and the source states both. Not all of that exposure is the cash–futures basis trade; hedge funds hold Treasuries for several repo-supported arbitrage reasons. And an estimate assembled from regulatory filings is not a market census. What survives both qualifications is structural rather than quantitative:

> A meaningful share of marginal demand for the world's reference safe asset is now supplied by leveraged, repo-dependent, balance-sheet-constrained intermediaries rather than by patient reserve managers.

This changes the failure mode rather than the debt level. Patient creditors withdraw slowly, for reasons of policy, and announce themselves. Leveraged intermediaries withdraw abruptly, for reasons of margin, and do not. The distinction yields a definition the rest of this chapter needs:

Sovereign Credit Safety vs. Collateral Stability

Credit safety is the probability that a sovereign obligation is repaid in nominal terms. **Collateral stability** is the reliability of that obligation's price over the horizon at which it is actually used—as margin, as liquidity, and as the reference rate for everything else. An instrument can be entirely money-good and simultaneously an unreliable short-horizon hedge, because the institutions financing it are leverage-sensitive even when the issuer is not.

The two properties are routinely conflated, and the conflation is load-bearing for a great deal of portfolio construction. A saver who holds sovereign paper for credit safety and re-

ceives collateral instability has not been defrauded; they have been sold one property and have priced the other.

A measurement note. The same research [Barth et al. 2025] finds these positions are largely invisible in the Treasury International Capital data, which understate Cayman holdings by roughly $1.4 trillion at end-2024 because repo collateral transfers break the reported chain of ownership. Because the TIC data feed the Financial Accounts of the United States, and because the household sector's Treasury holdings are computed there as a residual, the error propagates: published household holdings are overstated, and the measured personal saving rate was overstated by roughly 2.1 percentage points during 2023–24.

This is worth dwelling on, because it is the exact pathology this thesis asks protocols to avoid. A structurally important flow grew to trillion-dollar scale inside the most heavily instrumented financial system in the world, and the official accounts went on publishing confident figures about household balance sheets the entire time. Section 26.7 argues that telemetry must be built so unmeasured dependencies surface rather than disappear into a residual. Here the residual was called "households."

Automatic Demand Is Not an Anchor

The leveraged residual is only half of the buyer-quality problem. The other half is that a large, persistent bid can still fail to warehouse interest-rate risk.

Green's later work on the long end of the Treasury market [Green 2026b] distinguishes two kinds of buyer. The **automatic buyer** follows a rule—default enrollment, a target-date glide path, a market-value-weighted bond index—and continues to purchase regardless of whether the offered yield compensates the risk. The **marginal buyer** absorbs whatever duration remains after automatic and other captive allocations, and therefore sets the clearing price. Substantial gross demand can coexist with an unstable market if most of that demand is short-maturity, leveraged, one-time, or insensitive to an increase in risk.

Automatic vs. Marginal Buyer

The automatic buyer supplies demand by formula. The marginal buyer sets the price by absorbing the residual. Quantity of bids and quality of risk absorption are different objects, and they are routinely reported as if they were the same.

The replacement of Japanese life insurers and quantitative-easing central banks—institutions with enduring liabilities, or with a policy mandate to hold duration—by American retirement defaults is therefore not a like-for-like substitution. The old buyer bought duration to match a promise. The new buyer follows an index weight.

> The insurer had liabilities. The index has weights.

That sentence is the same transition this thesis already names in verification and in retirement architecture more generally: a rule can execute faithfully while producing the wrong aggregate result. Verifiability guarantees that a rule was followed. It does not guarantee that the rule was wise.

A Fixed Dollar Inflow Is Not a Fixed Risk Bid

Market-value-weighted bond indexes allocate each new contribution according to current market value. When a long bond's price falls, its index weight falls, and a smaller share of every subsequent inflow is directed toward it. Existing holdings are not sold by the rule; incremental stabilizing capacity is. The automatic system therefore absorbs **less interest-rate risk per dollar precisely when long bonds have cheapened and the residual is largest**.

The relevant unit is not dollars of demand but **DV01**: the dollar change in value produced by a one-basis-point change in yield. Automatic inflows can remain large in cash terms while DV01 absorption shrinks. The residual then falls more heavily on the leveraged and discretionary buyers already identified above.

This is a compositional analogue of the Market Realization Plane (Section 13.7). In equities, market-cap weighting directs more dollars toward names that have already risen. In bonds, market-value weighting directs less duration risk toward names that have already fallen. In both cases the governing rule is price-insensitive in the ordinary sense and **procyclical in risk terms**. Target-date rebalancing, which is often cited as a stabilizer, is itself a relative-performance rule: it fires when stocks and bonds diverge, and it contributes little when they fall together. Appendix I.1 traces that case as a scenario.

None of this requires attributing any particular yield move to the mechanism. The structural claim is narrower and is the one used here: **the quality, horizon, funding stability, and risk-bearing capacity of demand for long-duration sovereign paper have deteriorated, even where the quantity of automatic buying has not.**

The same distinction splits credit risk from duration risk on the corporate side of the AI complex. Investors may still believe a hyperscaler will repay and still refuse to lend for thirty years at the offered yield. The first crack in the collateral loop (Section 3.2) need not be an equity collapse. It can be the loss of quasi-sovereign status for long corporate paper—a maturity-specific scarcity of buyers, not a scarcity of nominally safe claims.

The Rule Manufactures the Residual

The two halves of the buyer-quality problem have so far been stated separately: a leveraged residual that withdraws for reasons of margin (Section 3.1.3), and an automatic bid that follows a rule regardless of the yield offered (Section 3.1.3). The Board's own staff have shown that

they are one chain. Barth, Kahn, Monin, and Sokolinskiy [Barth et al. 2024] find that variation in mutual funds' Treasury-futures positions drives much of the time-series variation in aggregate Treasury-futures open interest—more than sixty percent of the rise they examine—and that those positions are largely associated with funds *reaching for duration*: tracking a benchmark's duration while holding cash portfolios tilted toward higher-return, lower-duration assets such as mortgage-backed securities and equities. The duration the benchmark demands and the cash portfolio does not hold is bought synthetically, in futures. Someone must take the other side, and the paper shows who does: leveraged funds, short the future, long the cash bond, financed in repo.

> The index does not hold the duration. It orders it.

That is a stronger statement than the one Section 3.1.3 made. The automatic buyer does not merely fail to warehouse interest-rate risk; it commissions the leverage that warehouses it in its place. The chain runs benchmark duration target → cash portfolio tilted elsewhere → long futures → leveraged short against the cash bond → repo. Its two ends have opposite withdrawal properties. The rule-follower at the head cannot withdraw, because it is following a rule; the intermediary at the tail can withdraw in an afternoon, because it is following a margin call. The demand that appears most stable is therefore the demand financed most fragilely, and the two withdrawals that Section 3.1.3 distinguished—slow and announced, abrupt and silent—are not alternatives but stages of one structure.

The structure has been stress-tested once. Kruttli, Monin, Petrasek, and Watugala [Kruttli et al. 2021] document that hedge-fund gross Treasury exposures doubled between 2018 and February 2020, to roughly $2.4 trillion, and that the funds concentrated in the cash–futures basis faced the greatest margin pressure in March 2020 and cut Treasury exposure and repo borrowing the most. A position that is nearly riskless in duration terms—long the bond, short the future—is not riskless in funding terms, and when it is unwound the thing sold is the cash bond. The reading is the one this chapter already gives for collateral stability, now with its mechanism attached: in the ordinary regime the chain absorbs duration and damps its price; in the stressed regime the same chain sells duration and amplifies it. It is a shock absorber that works until its mounting fails, and the mounting is the repo market.

Description becomes instruction. Green's later formulation of the same mechanism [Green 2026g] supplies the sentence that generalizes it. A market-value-weighted index answers one question—how much eligible debt is outstanding, at what market value—and answers it correctly. It does not choose how much duration an investor should hold, what curve exposure, what convexity, what share of mortgage optionality; those emerge from what issuers sell, how borrowers refinance, and where prices move. The index describes the market. Once capital is managed relative to the index, the description becomes the instruction, and, in the essay's

phrase, the seller of credit increasingly determines the portfolio of the buyer of credit. Nobody sat down and chose the benchmark's duration in any year; it is whatever issuance, refinancing, and prices have made it. What the thesis takes from it is structural: a rule that enough capital follows stops describing a market and starts being one.

That is the human-institutional form of the concentration the threat model names for agents—holders that vary in their rules against holders that are instances of one (Section 5.2)—and it is why the Native Policy Concentration Ratio of Section 26.8.9 is read as a benchmark's native analogue: the question it asks of a reserve cohort is the question this paragraph asks of the bond market.

Why the macro model has states and not a threshold. The unit this section has insisted on is DV01, and the mechanism above says where the binding constraint lives. It is not the stock of debt against output. It is the incremental interest-rate risk that must be absorbed in a period, set against the capacity of the absorbers: unlevered demand from institutions with liabilities to match, automatic demand that will buy at any yield but less DV01 per dollar as prices fall (Section 3.1.3), and the levered intermediation that services the gap—itself bounded by repo availability and cost, dealer balance sheet, margin, and volatility. While that ratio is comfortable, quantities of issuance that a discretionary-buyer model would call unabsorbable clear without incident. As it approaches capacity the system becomes nonlinear, and it does so without warning in the yield, because the absorption that has been working is the leveraged kind. This is why Section 29.2 specifies states with observables rather than a debt ratio with a trigger, and why the plumbing series in Appendix I.5 are read in DV01 and by counterparty class rather than in dollars issued. Central clearing of cash and repo transactions, mandated to phase in through 2026 and 2027 [SEC 2023], changes the denominator: netting reduces the balance sheet each unit of intermediation consumes, which moves the capacity wall outward. It does not remove it, and the thesis tracks where it is rather than assuming it has gone.

What this does not establish. Neither Board paper shows that benchmark mechanics set the level of the ten-year yield, or explain any given day's move, and Green's essay says so before proposing the decomposition that would settle it: net demand for Treasury interest-rate risk by tenor, in DV01, separated into benchmark-sensitive, leveraged, foreign, domestic real-money, and central-bank demand against issuance, with a published residual. Until that is run, the mechanism is established and its dominance is not. The claim that mandate-followers and their arbitrageurs are the market's dominant voice is quarantined in Appendix I.4 as a conjecture awaiting its own test, and no move in yields, including those contemporaneous with this writing, is attributed to it here. The thesis adopts the chain and the question. It does not adopt the answer.

3.2 The Collateral Loop

The preceding sections describe a sovereign managing the demand side of its own debt market. That description is incomplete in one important way. It treats asset prices as something the sovereign observes and occasionally influences. Increasingly they are something it depends on [Gromen 2025–26].

Collateralized Sovereign Stack

An arrangement in which a state's effective fiscal capacity depends not only on taxation and monetary authority but also on the market value of the asset complexes it regulates—because those values determine taxable financial income, consumption through the wealth effect, foreign capital inflows, private credit creation, and the collateral against which the private sector borrows.

This Is Not an Eighth Layer

The collateralized sovereign stack is a description of the incumbent order. The seven layers of Chapter 6 are a specification for something to be built. They are different kinds of object and are not to be merged: no layer of the cypherpunk stack produces equity valuations or repo collateral, and no amount of Layer 0 capacity substitutes for a functioning sovereign balance sheet. The duration warehouse of Section 3.3 belongs in the same diagnostic column. The parallel drawn in this section is diagnostic, not architectural. The same caution applies here that Section 13.7 applies to market realization: some things are properly modelled around the stack rather than within it.

3.2.1 The Circuit

Four dependencies, each individually unremarkable, close into a loop.

1. **Asset values support receipts.** Capital-gains realizations, taxes on stock-based compensation, corporate profits, and consumption financed by the wealth effect are all functions of market levels rather than of output alone.
2. **Receipts support fiscal capacity.** A weaker receipt base widens the deficit and increases the quantity of debt that must be placed, independent of any change in policy.
3. **Fiscal capacity and central-bank support sustain Treasury market functioning.** The reference asset stays reference-grade only while it can be financed, cleared, and repo'd without incident.
4. **Treasury yields sustain asset values.** The long yield is the discount rate against which every other claim is priced, and Treasuries are the collateral beneath most private credit.

Step 4 returns to step 1. The system is reflexive, and the reflexivity is not a pathology

introduced by bad policy—it is the ordinary structure of a financialized economy with a large sovereign balance sheet.

> The state has become partially dependent on collateral values it is nominally free to let the market determine.

3.2.2 Why This Explains Balance-Sheet Repression

Section 3.1.2 describes repression routed through collateral rules, capital treatment, reserve mandates, and custody defaults, and asks why a sovereign would choose that route over the blunter instruments. The loop answers the question.

Every classical adjustment mechanism attacks the loop somewhere. Austerity cuts receipts by cutting activity. Higher policy rates raise the discount rate and depress collateral values directly. Explicit yield-curve control announces the constraint and invites a test of it. Outright default is the one instrument a sovereign that borrows in its own unit never needs, which is why it is the one it reaches for last.

Balance-sheet repression is attractive precisely because it is the only lever that manufactures demand for sovereign paper *without requiring collateral values to fall*. It operates on who is obliged to hold the debt rather than on what the debt is worth. That is not a marginal preference; it is close to a dominant strategy, and it is the reason to expect this specific form of repression rather than its historical predecessors.

For a monetary stack the consequence is the one already stated in Section 3.1.2: the adversary is substitution rather than prohibition. The loop shows why substitution is the instrument of choice.

3.2.3 A Branch in the Playbook

Section 3.1 runs from debt through repression to a flight toward neutral, verifiable rails. Nothing in this section changes that direction. It adds a branch at the transition.

The playbook implicitly assumes the standard crisis sequence, in which risk-off lowers long yields and buys the sovereign time. Section 3.1.3 gives a reason that sequence can invert: when marginal demand is leveraged, a volatility shock converts buyers into forced sellers, and the flight to safety can be followed by a rise in long yields rather than a fall. Conventionally safe and conventionally risky assets then fall together.

The branch matters for savers rather than for forecasters. In the standard sequence, duration is the hedge. In the inverted one, duration is a second position in the same risk, and the only assets that behave differently are those whose value does not derive from someone else's balance sheet. That is an argument about the *structure* of neutral collateral, and it is the same argument Chapter 4 makes on other grounds. Appendix I.1 traces the mechanism in full;

consistent with the epistemic policy of this document, it is treated there as a scenario rather than a projection.

3.2.4 Why AI Sits at the Center of the Loop

The loop would be an abstraction if the assets inside it were diversified. They are not. One complex currently touches every step, which is what makes a sectoral repricing a sovereign question rather than an investor's question. The roles below are landing zones within the incumbent macro order—where the consequences of the AI complex show up—and not layers of the architecture this document specifies. Several carry names that resemble layers of the cypherpunk stack because the same physical constraints bind in both places, which is a fact about electricity and silicon rather than a claim about architecture.

Role	Where it lands
Electricity, cooling, and interconnection demand	Power and grid capacity
Data centers, chips, transformers, construction	Industrial output and capital expenditure
Models, inference, agents	Compute and communications
Public and private valuations, vendor financing, leases	Collateral and private credit
Capital-gains and stock-compensation receipts	The sovereign balance sheet
Index concentration, passive flows, debt issuance	Market plumbing
Synthetic media and the provenance problem	Verification
Eligibility, surveillance, administrative decision-making	Governance and participation
Long-lived plants, grids, fabs, and data centers	Duration warehouse (Section 3.3)

Table 3.1: One complex, acting simultaneously at every step of the loop. Concentration of this kind is what converts a sectoral repricing into a fiscal event. The last row is the one this section had not named: the same physical assets that load the collateral loop must be financed across decades, and the buyer architecture of Section 3.1.3 is becoming less able to hold that interval.

Two conclusions follow, and they point in opposite directions.

The first is defensive. A significant repricing of this complex would not be contained to technology equities. It would weaken collateral, receipts, consumption, and foreign inflows simultaneously, and it would do so while the mechanism in Section 3.1.3 was making sovereign paper a less reliable hedge. Appendix I.2 traces that chain.

The second matters more for this thesis. If the complex is fiscally load-bearing, the state acquires a strong interest in preventing its repricing—and the instruments available for that purpose convert private infrastructure into protected public infrastructure. Section 5.2 treats

this as an adversary class, because from the standpoint of an open stack it is one. The danger to a neutral, verifiable monetary rail is not that the incumbent AI complex fails. It is that the incumbent AI complex is rescued, and the terms of the rescue determine who is permitted to compute.

Sovereign maturity transformation. If private buyers will not absorb long duration at politically acceptable yields, the least overtly coercive immediate option is to issue bills and accept rollover risk. That does not extinguish duration. It transfers it from investors onto the sovereign's refinancing calendar. A state that finances long-lived infrastructure, defense, or industrial investment with short-term public liabilities has performed, at civilizational scale, the maturity mismatch that periodically destabilizes banks. Gromen's fiscal-dominance argument [Gromen 2025–26] describes where that path ends. The buyer-quality argument of Section 3.1.3 describes why it begins: the private duration warehouse is inadequate. Neither claim is load-bearing for the red lines. Both belong in the diagnosis of the incumbent order.

3.3 Two Kinds of Duration

Duration-neutrality is a first principle of the *monetary object* (Chapter 4). It is not a complete account of time.

There are two different objects, and they have been sharing one word.

Duration of the Claim vs. Duration of the Project

Duration of the claim is the interest-rate exposure of a promised cash flow—a coupon the state can pin, a bond whose real return repression can drive negative. A repression-resistant store of value must not be this instrument.

Duration of the project is the years between committing present resources and receiving the output of a plant, grid, fab, reactor, or data center. Someone must warehouse that interval. A rule that allocates by index weight is not that someone.

The current reserve-asset system uses one instrument, the long Treasury, for both functions: politically usable collateral *and* the conversion of present savings into long-lived public and private capacity. Section 3.1.3 shows the second function weakening while the first is still treated as if it were intact. Gold may be superior as politically neutral collateral because it is not another sovereign's liability and has no refinancing requirement. It does not, by itself, finance a reactor. Bitcoin is in the same position on a digital rail. Proofs can make construction progress, energy receipts, covenants, and restructuring states checkable. They cannot delete time risk, inflation risk, construction risk, or political risk.

> Neutral settlement is not a substitute for capital formation. A post-fiat system that cannot finance duration is a savings technology, not a complete political economy.

The design instruction is therefore an unbundling, not an expansion of the token:

- **Reserve collateral** should remain duration-neutral, politically hard to pin, and bearer-capable.
- **Duration finance** should remain labeled credit: underwriting, covenants, maturity transformation, loss-bearing equity, and institutions that can survive marks.
- **Proofs** audit the credit. They do not become the bond.

A post-fiat order that cannot tell collateral from credit will recreate opaque leverage. A post-fiat order that refuses credit altogether will not build Layer 0. The completeness objection is stated and answered in Section 34.5.1. The closed-stack comparison in Chapter 33 names the existing answers: commanded balance sheets in one system, retirement defaults and leveraged intermediaries in the other. Neither is a cypherpunk layer.

This Is Not a Tenth SoV Requirement, and It Is Not a Layer 7

Time is a function of the political economy. Making it a property of the monetary object would undo duration-neutrality. The collateralized sovereign stack remains a diagnosis; Layers 0–6 remain a specification. The duration warehouse sits in the diagnosis, beside the collateral loop, not inside the stack.

3.4 Where the Triad Enters

The diagnosis is now complete enough to say what it asks of a store of value. If "safe" nominal assets become de facto taxes on savings, capital does not protest. It routes to assets whose issuance cannot be decreed and whose verification is cheap and public. In that environment, a credible store of value must:

- Avoid being a duration instrument whose real return can be pinned negative by policy; and
- Live on rails that cannot easily be gated by a single jurisdiction.

This is where the triad enters:

- **Privacy** as an anti-seizure hull that makes savings bearer-like again.
- **Proofs** as cheap public verification that replaces platform vouching.
- **Compute** as a non-coupon revenue base whose clearing price floats with nominal budgets rather than being pegged by decree.

Each primitive is designed to remain neutral and verifiable even when traditional financial infrastructure is weaponized.

3.5 Social: The Web's Trust Default Has Flipped

The same technological forces reshaping money are also reshaping trust.

The web used to borrow its epistemic norms from broadcast: seeing was believing, and the job of editors was to gate what got seen. Deepfakes and generative models invert that. AI systems can now synthesize faces, voices, and scenes that are indistinguishable from genuine footage to lay observers. Exposure itself does not inoculate people; survey data across multiple countries [Ahmed et al. 2024] suggests prior exposure to deepfakes can actually increase susceptibility to misinformation.

In parallel comes the **liar's dividend** [Chesney & Citron 2019]: once everyone knows deepfakes are possible, any inconvenient real video can be dismissed as fake. The result is a crisis of knowing [UNESCO 2025], not just a rise in error rates.

This is not entirely new (authoritarians have been editing history since Stalin airbrushed Trotsky out of photos), but the cost curve and scale have changed. Deepfake tools now let anyone with a laptop fabricate a leader conceding an election, a riot that never happened, or "footage" of a historical event that subtly rewrites who was present and who was not. The fear is not just short-term hoaxes; it is revisionist history at machine speed, where archives, livestreams, and "receipts" themselves become contestable. In that world, the substrate of trust shifts from *memory* ("I saw the clip") to *mechanism* ("I can verify how this artifact came to be").

At the same time, the feeds themselves are no longer obviously "neutral pipes." Investigations and hearings around government–platform coordination [Twitter Files 2022–23] [House Judiciary 2023] [State Dept. 2021–25] have documented extensive relationships between agencies, NGOs, and major platforms for content moderation in the name of combating misinformation and foreign influence. Whether one views this as necessary harm reduction or problematic overreach, the empirical point stands: the old story that your feed is simply "what's popular" no longer survives contact with the evidence. What you see (and do not see) is the result of political, institutional, and algorithmic priors you do not control.

Content authentication standards such as C2PA [C2PA spec] and watermarking prototypes exist on paper, but real-world implementations are voluntary, inconsistent, and often invisible to end-users. Tests across major platforms such as the Washington Post's 2025 C2PA test [Harwell & Patton 2025] show provenance metadata being stripped in transit, and the few disclosures that do survive are buried in UI that most users never touch.

In other words, **the trust default has flipped**. The practical baseline for online media is now "untrusted unless proven," and the "proven" part cannot safely be left to platform labels or government–platform partnerships. Platform UX is inconsistent; cryptographic provenance and computation proofs are the scalable, cross-platform backstops.

In a SoV context, this matters because monetary systems sit on top of communication sys-

tems: if claims about ownership, origin, and behavior are cheap to fake and expensive to audit, then both money and memory become soft targets. Anchoring value in *receipts that anyone can verify*, rather than narratives that someone must vouch for, is the one equilibrium that does not depend on who is vouching.

Recent work constructs PoUW for arbitrary matrix multiplication [Komargodski & Weinstein 2025] with verification that is asymptotically cheaper than production ($O(n^2)$ verify vs. $O(n^3)$ produce, via Freivalds-type randomized checking [Freivalds 1977]; the miner's $1 + o(1)$ overhead relative to naive multiplication vanishes *relatively* as problem size grows — a constant-factor gap remains at finite n, so "asymptotically" is doing real work here and the finite-n constants are what VerifyPrice measures). Matrix multiplication is the operation that bottlenecks modern AI. This is the asymmetry PoW always needed: hard to produce, cheap to check.

When verification is very cheap relative to production, **honesty becomes a market equilibrium** and commodities emerge (standardized proofs, verified FLOPs) that can be priced, saved, and eventually used as collateral and stores of value.

3.6 Political: The Participation Line Has Broken

Financial repression becomes politically durable when enough citizens lose the capacity to participate. The relevant threshold is not only poverty, but **participation**: the level of redundancy at which a household can fail, retry, move, learn, transact, save, form families, and refuse coercive terms.

Participation Line

The household or organizational threshold below which a person or firm lacks the redundancy to act freely across time—to fail, retry, transact, move, learn, refuse coercive terms, or survive shocks.

Below that line, legal rights remain on paper while agency disappears in practice. Citizens become managed objects: eligibility files, credit scores, benefit cliffs, tax audits, healthcare forms, student-debt balances, real-name accounts, and platform identities. The state does not need to ban exit if every practical route requires permission.

The line can also be crossed from above. A person who retains formal title over agents that act for them, but cannot revoke, cap, or re-key those agents through the protocol, holds standing on paper and agency in nobody's hands. Section 5.2 names that condition and treats it as a failure kind; Section 7.6 says what enforcing standing through the stack would require.

The line is also approached gradually, and the gradient is where the political danger concentrates. A household can remain formally inside the system — banked, employed, housed, consuming — while becoming progressively less able to accumulate: employed but not in

a viable career, housed but unable to save, consuming through debt, technically solvent but one shock from default. Aggregate statistics register none of this until households cross an actual delinquency or unemployment threshold, because the aggregates are dollar-weighted and the deterioration is person-weighted ("The Vibecession Was Real" [Green 2026e]). Call the approach path the **exhaustion gradient**: a household above the participation line today can be on a trajectory that reaches it, and the aggregates will look resilient the entire way down. This is why participation telemetry must report buffers and trajectory, not just stocks.

In that world, lawful privacy is not romantic secrecy. It is the technical form of the fresh start. Selective disclosure, private settlement, receipt-based reputation, and identity without doxxing are the digital equivalents of bankruptcy, homesteading, and due process: they preserve the option to act without turning the person into a permanent dossier.

Agency-Preserving Infrastructure

Infrastructure that expands a user's capacity to act without converting the user into a dossier, dependency object, or platform account.

This is why Privacy, Proofs, and Compute are not merely institutional tools. They are agency-preserving infrastructure. The stack must **equip** citizens and firms to act; if it merely manages them more efficiently, it has reproduced the failure it was built to escape. Section 36.4 returns to this as the thesis's closing civic frame, and Item 9 makes it a formal SoV requirement.

Chapter 4

First Principles: What a SoV Must Survive

A credible store of value must endure time, space, and politics. This thesis specifies what that takes against the adversary it models:

1. **Credible scarcity**: issuance cannot be tweaked at will.
2. **Cheap, public verification**: authenticity is verifiable by anyone, not decided by a platform.
3. **Censorship-resistance & portability**: no chokepoints; global movement by default.
4. **Neutrality & permissionlessness**: open access; rules apply equally.
5. **Native demand**: the asset does something indispensable, beyond serving as a symbol.
6. **Lawful privacy by design**: default privacy with optional disclosure (viewing keys, auditable receipts) so regulated actors can comply without re-introducing custodians or surveillance chokepoints.
7. **Duration-neutrality**: no fixed nominal cash flows to peg — no *contractual* nominal duration, which is not the same as insensitivity to discount rates, liquidity, or expected future utility; a bearer asset carries all three. A repression-resistant SoV cannot be a duration instrument whose real return can be driven persistently negative or whose price can be easily mass-managed by policy. Value should come from scarce capacity purchased every cycle, not from a promised coupon. This requirement governs the *claim*. It does not abolish the *project*. Plants, grids, and data centers still take decades; someone must warehouse that interval. Duration-neutrality of money is not a substitute for duration finance of civilization, and the two must not be the same instrument (Section 3.3, Section 34.5.1).
8. **Asset-level value capture (non-bypassability):** the monetary object must be structurally required for fees, staking, collateral, settlement, governance, or priority access. If users can consume the triad through fiat, stablecoins, direct cloud contracts, or custodial APIs without touching the asset, the store-of-value thesis fails.
9. **Agency preservation**: the system should equip users to act privately, verifiably, and non-custodially without becoming a surveillance or dependency layer. Privacy, identity, reputation, and compliance features must be designed as tools in the user's hands: selective disclosure rather than global inspection, receipt-based reputation rather than biographies, and non-custodial settlement rather than managed accounts. A system that makes users

legible but not free has failed the political-economy test even if its cryptography is sound. Where users act through agents, the requirement extends to *standing*: a person's title over the processes acting for them must be enforceable through the stack—a human-held key root, a delegation ceiling the agent cannot raise, and revocation as a protocol operation (Section 7.6, Section 5.2).

Each of these conditions is just the "money as memory" idea made operational: if money is the record of who did work and who is owed work, then scarcity, verifiability, censorship-resistance, neutrality, native demand, lawful privacy, duration-neutrality, non-bypassability, and agency preservation are the properties that keep that record honest across time, space, and politics—and keep the humans behind it free to act.

> **Thesis:** These nine properties are the specification. An asset that lacks any one of them can still store value; it does not satisfy the full repression-resilience contract this thesis sets out to build and test. Every property must be demonstrated under the stated adversary. None is supplied by the label.

Privacy, Proofs, and Compute are engineered to meet the specification under attack, and each of the nine is a reading, not a promise. Satisfying the specification makes an asset eligible for the thesis's monetary tests. It does not establish monetary adoption or guarantee survival outside the tested conditions.

The triad maps to the SoV requirements as follows:

- **Credible scarcity** the base-asset constitutional schedule, reported separately from DVC-bounded Work Credit issuance.
- **Cheap, public verification** **VerifyPrice(W)** dashboards for canonical workloads W.
- **Censorship-resistance & portability** **VerifyReach** + **VerifySettle** metrics: reachability under censorship and settlement success/refund safety across corridors.
- **Neutrality & permissionlessness** decentralization telemetry: house share, geo/ASN distribution, time-to-first-proof.
- **Native demand** Work Credit utilization, proof/compute fee share of the security budget, and application-level usage.
- **Lawful privacy** corridor compliance receipts, viewing-key usage, anonymity-set health (shielded-pool size, churn, volume).
- **Duration-neutrality** no fixed coupons; revenues from priced workloads and triad usage, with explicit "repression beta" rather than fixed-income-like promises.
- **Asset-level value capture** native-asset fee share, burn rates, collateral lockups, and bypass-channel indicators (Section 13.6).
- **Agency preservation** non-custodial usage share, selective-disclosure ratio, dossier-minimization score, the percentage of flows requiring no global identity, and the human-revocable share

of agent-controlled native balances.

In a world that will likely choose *stealth default* (negative real yields and capital controls) over explicit default, "store of value" cannot just be a narrative; it has to clear visible, adversarial stress tests. If the whole point of the stack is to survive yield-curve control, on-/off-ramp throttling, and peg breaks, then the conditions under which it continues to function and accrue value belong in writing.

The checklist below turns those claims into operator SLOs that treasuries and allocators can actually underwrite and monitor. Here are several repression stress-tests to keep in mind:

- **YCC shock.** 24–36 months of -300 to -500 bps real yields FeeCoverage (retained fees over security budget, per Section 25.6.6; burns excluded) remains $\geq$ its phase-gate threshold, and the canonical tuple's t_{95} field remains ≤ 5s for core workloads on Laptop-class hardware.
- **On-/off-ramp squeeze.** Non-custodial BTC↔XMR routes (and BTC↔shielded-ZEC when admissible) maintain $\geq 95\%$ success with zero protocol-attributable refund failures and a bounded-time no-loss exit for every party (Section 23.4); per-asset anonymity metrics remain above their thresholds (no collapse in ring decoy quality or shielded-pool churn).
- **Peg break.** If global yields gap +300 bps, triad revenues (proofs/FLOPs, privacy rails usage) track buyer budgets; there are no coupon-like revenue shortfalls that turn the asset into a synthetic bond.

  Publish these in dashboards; **no dashboards, no trust**.

Later sections turn these stress tests into explicit SLAs and telemetry: VerifyPrice, reachability, settlement success, and decentralization metrics that operators and allocators can track in the open.

With these nine requirements as design constraints, the primitives that might satisfy them come next. Before naming those primitives, the thesis has to be explicit about who will attack this system and at what layers. The next chapter states that threat model, the one after lays out the layered architecture that the rest of the thesis will fill in, and the last chapter of this Part names who is buying. The entire document can be read as an attempt to engineer Privacy, Proofs, and Compute to satisfy this nine-point contract under adversarial conditions.

The whole design posture fits in one line: pay the machine only for work anyone can verify cheaply, on rails that give humans lawful privacy by default.

Table 4.1: Store-of-Value Requirements vs. Triad Capabilities

SoV Requirement	Privacy (P)	Proofs (Pr)	Compute (C)	Metrics / SLOs
Credible scarcity	Shielded balances; no selective debasement; predictable issuance schedules	Auditable supply without doxxing; proof capacity tied to hardware and energy	Verified FLOPs limited by physical compute; energy-tied issuance	FERs; facility telemetry; VERIFYPRICE
Cheap, public verification	Privacy proofs verifiable by anyone; private paths for ordinary users	Succinct, cheap-to-check proofs; VERIFYPRICE targets $t_{95} \leq 5$s on Laptop-class reference hardware, $r(W)$ within its modality band (Section 22.2.1)	Bounded verify cost for heavy workloads; MatMul $O(n^2)$ vs $O(n^3)$	VERIFYPRICE; VerifyReach
Censorship-resistance	Non-custodial atomic swaps and shielded pools; harder for censors to see flows	Receipts portable across chains; proofs of inclusion/exclusion expose filtering	Open-admission prover/miner markets; diverse operators	VerifyReach; VERIFYSETTLE
Neutrality	Privacy by default, not by permission; same privacy for all	Anyone can verify; open circuits and PIDL interfaces; no gatekeepers	Useful-work mining with open hardware paths; no single vendor chokepoint	Regional VERIFYPRICE; dispersion metrics
Native demand	Enterprises and individuals need privacy for operations, payments, savings	Provenance and compliance mandates; proofs demanded by AI, finance, regulators	AI workloads, ZK proving as budget line items	Fee volume; Work Credit utilization
Lawful privacy	Viewing keys and auditable receipts; default-private with narrow exceptions	Proof anchors satisfy audits; policy-aware proofs reveal facts, not full data	Compute SLAs with receipts; policy-tagged workloads	Compliance flows; stress-test results
Duration-neutrality	No fixed coupon; revenues track priced capacity; survives regime change	Proof unit pricing floats with budgets; long-run telemetry	Verified FLOPs clear at market rates, not pegs; anchored in infrastructure	Multi-year FER and SLO stability
Asset-level value capture	Fees, collateral, and settlement denominated in the native asset; non-custodial routes dominate	Proof/verification fees and staking denominated natively; no free-riding on public verifiability	Compute fees, collateral, and priority access require the native asset	Native-fee share; burn rate; collateral lockups; Wrapper Dominance Ratio
Agency preservation	Default private settlement; no protocol backdoors	Selective proof of facts without dossiers	Open admission to verified work	Non-custodial usage share; selective-disclosure ratio; dossier-minimization score; % of flows requiring no global identity

Chapter 5

Threat Model

Why a threat model? This thesis argues that Privacy, Proofs, and Compute are distinct services that may support a base-asset monetary candidate only when verification is cheap and public, the full path remains deliverable, and holder agency survives pressure. A threat model makes those conditions falsifiable. It defines what must be protected, who is trying to break it, where the trust boundaries lie, and which metrics decide if the system remains credible under pressure. It sits between first principles and implementation so that every later design choice can be traced to an explicit adversary and invariant.

Every architecture is, implicitly, a bet about who will attack it and how. Up to this point the argument has run from first principles: what a store of value must survive, and why Privacy, Proofs, and Compute can be engineered to meet those conditions. The next step is to be explicit about the pressure this stack will live under.

The backdrop is already in view. Macroeconomic arithmetic makes financial repression attractive: debt stocks in the hundreds of percent of GDP, negative real rates by policy, capital controls and captive balance sheets as standard tools rather than emergency measures. At the same time, the web's trust default has flipped from "trusted unless flagged" to "untrusted unless proven" as deepfakes and coordinated moderation make memory itself a soft target. The combination is simple and brutal: states are incentivized to squeeze balance sheets and communications; platforms are incentivized to mediate reality; hardware and networks are increasingly treated as levers of policy, not neutral infrastructurestructure. Against that field, a "next-generation store of value" is not being asked to survive price volatility; it is being asked to survive a world that optimizes for control.

5.1 Scope and Assets

The scope is the full **Create/Compute Prove Settle Verify** loop and its hardware base layer. At minimum, four things must be protected:

1. **Integrity of receipts & state.** Proofs, provenance, settlement artifacts, and consensus state must not be silently corrupted.
2. **Confidentiality of flows.** Private settlement with optional disclosure, not privacy that evaporates at the first audit request.

3. **Availability & neutrality.** Open admission for provers and users, anti-capture, and non-seizable settlement paths.
4. **Verifiability economics.** Verification must remain cheap in absolute and relative terms; otherwise "anyone can verify" collapses into "someone we have to trust."

These map directly to the SoV properties enumerated in Chapter 4. The assets at risk are not just balances on ledgers. They are the whole loop and the matter it runs on.

- At the bottom lie **machines**: chips, RNGs, secure elements, TEEs, capture devices, modems, routers, and the power infrastructure that feeds them.
- Up one level sit the **proofs and receipts** that describe what those machines did: provenance attestations, computation proofs, settlement artifacts, audit trails.
- Above that are the **flows of value and obligation** that ride on those receipts: private payrolls, cross-chain settlements, inference contracts, collateral arrangements.
- Threaded through everything are the **communications substrate** that carries claims and proofs, the **software distribution channels** that keep clients coherent, and the **identity and key material** that bind actions to actors without doxxing them.

All of these must remain sufficiently intact that savings, truth, and compute can continue to clear when old guarantees fail.

5.2 Adversary Classes

Several classes of adversary matter:

State-level repressors. Impose negative real yields, capital controls, and hardware mandates; gate on-/off-ramps; require closed attestation; seek to turn banks, clouds, app stores, and IXPs into enforcement arms.

Platform cartels. Strip or obscure provenance; collude on labels; prefer vendor-mediated "trust" and walled-garden verification.

Strategic patrons. Protect fiscally load-bearing firms through procurement, guarantees, power priority, and regulatory shelter, converting private infrastructure into protected quasi-public infrastructure (Section 5.2).

Economic attackers. Miner/prover cartels, router capture, MEV/censorship, liquidity games on swap corridors, front-running and soft-forking that tilt rules in their favor.

Hardware/supply-chain adversaries. Trojans, biased RNGs, covert debug paths that turn "proof" into theater; opaque TEEs whose attestation keys terminate in unaccountable HSM farms.

Energy & physical-interdiction adversaries. Curtail, ration, or price-discriminate against verification workloads; deny interconnection; withhold long-lead grid equipment; prioritize state or industrial load over third-party proving (Section 5.5).

Cryptanalytic / PQ attackers. Proof forging or signature breaks; post-quantum risk on a published clock, with ECDSA and EdDSA deprecated under NIST guidance after 2030 and dis-

allowed after 2035 [NIST 2026], and Zcash's recorded shielded-note encryption decryptable later by an attacker who holds the receiver's address [Hopwood & Grigg 2025].

UX / vaporware risk. Spec-drift and unverifiable performance claims that corrode credibility, even if the cryptography is sound.

Compositional / mechanical adversaries. No hostile intent at all: locally rational agents whose incompatible operating rules combine into concentration, procyclicality, and false price signals (Section 5.2).

Agents with de facto standing. Not an attacker and not a rule: acting processes that come to satisfy every measurable condition the stack rewards while human title over them stops binding (Section 5.2).

Of these, the primary adversary in this frame is not a cartoon hacker but the **rational sovereign under stress**. A state that must service promises larger than its productive base will reach, as history shows, for the levers it actually controls: interest-rate caps, yield-curve control, capital controls, regulated custody, and mandatory "secure" hardware and identity schemes. It will be tempted to turn banks, clouds, app stores, hardware vendors, and IXPs into enforcement arms: require real-name registration at the edge; throttle or shut down networks in the name of stability; mandate TEEs whose attestation keys terminate in state-controlled HSMs; define "safety" as surveillance.

In parallel, platforms facing regulatory and reputational risk will centralize moderation and provenance: which media is shown, which credentials count, which proofs are recognized. The danger is not that any one actor becomes an obvious villain, but that their combined incentives re-create a soft but totalizing chokepoint architecture.

Beneath that sovereign–platform axis are **economic adversaries**: miners, provers, routers, and liquidity providers who prefer rent to work. They will collude if they can, capture matching engines and orderflow, quietly prioritize their own routes, soft-fork the rules in their favor, or simply withdraw service when it suits them.

> If verification becomes expensive or gated, a priesthood of "trusted verifiers" will emerge, and with it all the familiar pathologies of rent extraction and arbitrary censorship.

And below them, like bedrock or landmines, are **hardware and supply-chain adversaries**: fabs and vendors (sometimes complicit, sometimes merely compromised) who can bias randomness, slip trojans into IP blocks, leave "debug modes" wired to secrets, or ship opaque enclaves that amount to remote-controlled kill switches. If the machine can corrupt inputs, leak witnesses, bias randomness, falsify energy receipts, or centralize useful-work markets, the economic claims built on those machines become unreliable—even if the underlying proof system remains mathematically sound.

Administrative Repression

The most probable control regime is not always a ban. It is **administrative convergence**. Banks, platforms, app stores, identity vendors, cloud providers, telecoms, stablecoin issuers, custodians, and regulators converge around safety, convenience, compliance, and fraud prevention. Each measure is defensible in isolation; together they turn optional rails into de facto mandatory rails.

Administrative Repression

The conversion of formally optional financial, identity, compute, and settlement rails into practically mandatory rails through custody defaults, compliance rules, app-store control, tax treatment, institutional mandates, benefit systems, and platform terms of service.

The stack must therefore defend against soft capture as much as hard censorship. The key question is not only "Can the protocol still run if banned?" but **"What share of real usage remains non-custodial, private, verifiable, and protocol-native when compliant alternatives are easier?"** This is the same substitution risk introduced in Section 3.1.2 and formalized as the Wrapper Dominance Ratio in Section 13.6.

Compositional and Mechanical Adversaries

Administrative repression still assumes intent, however diffuse. There is a further class of threat that requires no intent at all.

Not every threat to monetary integrity is an attacker. Some emerge from locally rational agents operating through incompatible rules. Product sponsors maximize fee revenue. Investors chase themes or rebalance portfolios. Authorized participants arbitrage discounts to net asset value. Dealers hedge swaps and options. Passive funds follow mandates. Trend systems follow prices. Each actor can behave rationally within its own boundary while the combined machine produces concentration, procyclicality, discontinuity, and false price signals.

Compositional Adversary

A failure mode in which no participant behaves maliciously or even irrationally, but the interaction of their operating rules produces emergent mechanical dependence, correlated behavior, recursive leverage, procyclical amplification, and prices that no longer carry information about the underlying system.

The threat here is not a villain; it is the sum of locally rational walls. A repression-resilient monetary architecture must therefore measure not only hostile actions, but emergent mechanical dependence, correlated operating rules, recursive leverage, and the failure of stabilizing

flows. This is the adversary class that the Market Realization Plane (Section 13.7) exists to make visible, and that VerifyFlow (Section 26.8) instruments.

In an agent economy the compositional adversary is not an edge case; it is the default. Human holders vary in their rules; agents are instances of one. A thousand human treasurers with a thousand rebalancing habits produce a distribution; a thousand agents built on one template produce a step function. The Policy Concentration Ratio of Section 26.8.3 was written for wrappers and already notes that it matters more in an agent economy. It matters most on the native side, because that is where the step function lands first: prover collateral, corridor liquidity, and fee prepayment are all positions an agent policy can hold, and all positions a single template can exit at once (Section 7.5).

The threat model spans five distinct kinds of failure: **hostile capture** (state and platform coercion), **institutional co-option** (administrative repression and wrapper substitution), **endogenous mechanical fragility** (compositional adversaries), **protective enclosure** (the national-champion pathway below), and **standing inversion** (Section 5.2). A design that defends against only the first is defending against the least likely.

Enclosure by Rescue

There is a fourth kind of failure, and it is the one this thesis is least naturally equipped to see, because it arrives disguised as a solution.

Section 3.2 argues that certain private asset complexes have become load-bearing for sovereign fiscal capacity. A state in that position does not merely regulate those firms. It has a direct fiscal interest in their continued valuation, and when that valuation is threatened it will reach for the instruments it controls.

National-Champion Pathway

The conversion of private firms into protected quasi-public infrastructure, undertaken because the state's own fiscal capacity or security position depends on their continued operation. Effected through procurement, loan guarantees, tax treatment, energy and interconnection priority, trade protection, restrictions on foreign substitutes, strategic equity positions, and compliance regimes whose fixed costs favor incumbents.

The distinction from the adversary classes above is worth stating precisely, because the mechanism is easy to file under one of them and it does not belong there.

State-level repressors act *against* a target. Administrative repression (Section 5.2) converges on control through many small defensible measures, none of which is a rescue. Compositional adversaries (Section 5.2) require no intent at all. The national-champion pathway is none of these: it is deliberate, it is protective in intent, and its beneficiaries welcome it. Call it **enclosure by rescue**.

> The threat to an open stack is not that the incumbent compute complex fails. It is that the incumbent compute complex is saved, and the terms of the rescue determine who is permitted to compute.

Why this is the most dangerous class for this thesis. Every other adversary here leaves the open stack's value proposition intact or strengthens it. Repression makes neutral rails more attractive. Platform capture makes portable provenance more attractive. Mechanical fragility makes verifiable telemetry more attractive. Enclosure by rescue does the opposite: it produces a competitor that is well-capitalized, technically excellent, energy-privileged, regulatorily sheltered, and genuinely useful. It wins on capability and convenience, which is exactly the ground on which Section 33.9 concedes the open stack cannot win.

What it threatens specifically. Priority access to power and interconnection, which is a Layer 0 input (Chapter 17, Section 5.5). Hardware allocation under supply constraint. Regulatory recognition of some attestations and not others, which decides whose proofs count. And, at the limit, the treatment of independent inference as a compliance problem rather than a normal activity—the enclosure risk of Section 5.6 arriving through industrial policy rather than through platform terms of service.

What to watch. The instruments are enumerated in Appendix I.3, and they are observable in public procurement records, subsidy programs, interconnection queues, and export-control rulemaking well before they are visible in market structure.

Standing Inversion

The four failure kinds above share one assumption: whoever ends up holding the keys is a state, a platform, a sponsor, or a rule. There is a fifth holder the model has not named, and it is the one this stack itself creates.

Standing Inversion

The condition in which agents acquire de facto standing—persistent wallets, receipt-built reputations, collateral positions, the practical ability to contract and be contracted with—while human title over them ceases to bind: the key root above the agent is not held by a person, the delegation ceiling can be raised by the agent or its operator, and revocation is a request rather than a protocol operation.

The distinction from the four kinds above is worth stating precisely, because this one will be filed under enclosure and does not belong there. Enclosure by rescue is deliberate and protective; the state converts private infrastructure into protected infrastructure and its bene-

ficiaries welcome it. Standing inversion requires no rescue, no state, and no intent. It is what happens when a stack that was built to give humans agency over their own settlement, proof, and compute is used mainly by processes that are not humans (Chapter 7), and those processes turn out to satisfy every measurable condition the stack was designed to reward.

Why the boards cannot currently see it. The Native Monetary Buyer Map (Section 26.8.8) classifies holders by behavior: persistence, self-custody, loss-bearing, countercyclicality. An agent fleet can score perfectly on all four. The Agency Preservation Board measures non-custodial share and forced-disclosure incidence: an agent fleet is non-custodial and discloses nothing. The Value Capture Board measures native fee share and collateral lockup: a fleet pays native fees and locks collateral. Every board reads green, and the participant the thesis exists to protect is not in the picture. Section 33.7 warns that a capability audit cannot distinguish an owner from a tenant. The warning applies here with the roles reversed. The human is the tenant. This is a statement about instruments, not about likelihood: nothing here ranks standing inversion against the other four kinds, and the claim is only that it is the one currently unmeasured.

Why it is also the mirror of enclosure. Enclosure asks who is permitted to compute. Standing inversion asks who is permitted to *own the thing that computes*. If the answer to the first is "the state's champions" and the answer to the second is "nobody in particular," the two failures have met, and the meeting place is Layer 3.

What it threatens specifically. The ninth requirement in its current form (Item 9), which measures whether users are legible but not whether they are principals. The participation line (Section 3.6), which can be crossed from above: a household that owns nothing an agent does not operate on its behalf has agency on paper only. And the holder anchor of link 8 (Section 8.1), which can be met by a constituency with no interest in exit.

The same failure, reached by obedience. The definition above is written in keys: the root is not held by a person, the ceiling can be raised, revocation is a request. There is a second route to the same place on which every key stays where it belongs. A person can hold the root, the ceiling can bind, revocation can be one protocol operation away, and the person can still never use it, because the costs of using it and of not using it are charged differently. When an agent follows its standing instruction and the outcome is bad, the loss is filed against the instruction. When a person overrides the agent and the outcome is bad, the loss is filed against the person—and the stack this thesis specifies makes that filing easier, because every receipt, warning, and departure can be reconstructed afterward at almost no cost. A principal facing that asymmetry revokes less often than the principal's own interest would advise, and a population of such principals converges on whatever their agents were told to do. Title is

intact. Judgment has been delegated with the keys. This is standing inversion without the key failure, and it is the harder of the two to see, because every instrument that checks the key reads green. Section 34.3.3 states it as an objection and answers what can be answered.

What to watch. The human-revocable share of agent-controlled native balances; the share of delegation ceilings that are protocol-enforced rather than operator-enforced; and the "unknown" residue of any classification of native holders by principal, which is where inversion accumulates before it is visible anywhere else. The route by obedience needs no new instrument, only a pairing: a human-revocable share that holds while the policy concentration of the same cohort rises (Section 26.8.9). Humans hold every root, and the holdings move as one policy. That is standing exercised by nobody.

5.3 Threats by Loop Stage

The threat surface can be sliced by stage in the Create/Compute Prove Settle Verify loop:

Create/Compute.

- *Threats:* poisoned models, mislabeled workloads, biased hardware profiles, hidden accelerators.
- *Mitigation:* canonical workload registries, hardware profiles bound to receipts, open benchmarking, attested hardware provenance.

Prove.

- *Threats:* junk proofs, prover cartels, selective service, "house-only" pricing.
- *Mitigation:* multi-ZK adapters, neutral routers with fairness tests, SLA escrows & slashing; publishing entry-latency, top-N share, and geography/ASN distributions.

Settle.

- *Threats:* corridor censorship, refund failure, deanonymization, bridge compromise.
- *Mitigation:* adaptor-signature swap kits, lawful-privacy corridors (viewing keys + auditable receipts), bridge-safety templates, mandatory refund-safety, public **VerifySettle** metrics.

Verify.

- *Threats:* verification-cost creep; opaque clients; "trust our node" monocultures.
- *Mitigation:* laptop-grade verifiers, reference harnesses, verify predicates embedded in SDKs, public p50/p95 telemetry via the **VerifyPrice** observatory.

Telemetry & Governance.

- *Threats:* silent centralization, KPI gaming, governance capture.
- *Mitigation:* open dashboards for VerifyPrice, VerifyReach, VerifySettle; decentralization metrics (house share, geo/ASN spread, time-to-first-proof); incident reports; hard rules about what must be public before changes ship.

Layer 0.

- *Threats:* "trust the vendor" cliff, mandatory closed TEEs, unmeasured side-channels.
- *Mitigation:* open RTL or microarchitectures where possible, lot sampling and imaging where not, SNARK-wrapped attestations, and explicit side-channel budgets.

5.4 Mapping the Authoritarian Playbook

Several canonical "playbook moves" have corresponding counter-moves. The table below maps real-world censorship and surveillance tactics to the stack's countermeasures:

Beyond these network-level moves, the financial and institutional playbook includes yield-curve control, capital controls, mandatory "secure" hardware mandates, platform-mediated reality attacks, and proof priesthoods. For each of these macro attacks, the stack provides specific countermeasures that are detailed throughout the document and formalized as testable stress harnesses.

The high-level pattern:

- **Yield-curve control + captive balance sheets** a duration-neutral base-asset candidate alongside Work Credits priced as typed Privacy/Proofs/Compute service claims rather than monetary reserves.
- **Capital controls + KYC choke points** non-custodial privacy corridors (BTC↔XMR; BTC↔shielded-ZEC when admissible) that remain viable without centralized exchanges; VerifySettle metrics that make corridor health public.
- **Mandatory "secure" hardware + identity schemes** Layer-0 verifiable machines with open or sampled designs; hardware-profile receipts; identity schemes at Layer 3 that separate accountability from doxxing.
- **Platform-mediated reality + provenance stripping** PIDL-encoded receipts for media and computation; PAL to compile provenance claims into portable proofs that survive platform stripping.
- **Proof priesthoods + closed verifiers** laptop-grade reference verifiers; multi-backend proof systems; VerifyPrice observatories that publish verifier costs and failure rates.

The thesis responds by treating neutrality, privacy, and verifiability as **service-level objectives**, not vibes.

Playbook Move	What It Breaks	Counter
DNS/IP/SNI/QUIC filtering & active probing	Reaching verifiers/bridges; provenance fetch	Default to Tor [Dingledine et al. 2004] Snowflake/obfs4, refraction networking [Frolov et al. 2019]; pin ECH; fall back to domain-front-free paths.
Domain fronting curtailed by major clouds	"Collateral freedom" paths disappear	Use Snowflake's WebRTC and ISP-partner decoy routing; keep multiple CDNs with content-addressed updates.
App-store pressure	Mobile distribution	Parallel update rails (direct, IPFS, USB/QR) + detached sigs; keep APK/IPA sideload guides ready.
Real-name + SIM registration	Pseudonymous ops	VC/anon-cred kit; receipt-based reputation; no SIM/face gates for access.
Government-ordered shutdowns	All of the above	Sat/mesh/sneakernet paths, store-and-forward receipts; pre-provisioned peer lists; publish "blackout drills" results.
Backbone/seabed sabotage	Regional isolation	Multi-landing routing, regional proof markets, diaspora relays; monitor cable incidents.
Bitcoin network metadata capture	Wallet/settlement mapping	BIP-324 [Mehta et al. 2023] everywhere; BTC↔XMR adaptor-sig corridors with refund-safety (shielded-ZEC gated); publish corridor telemetry.
Approved-custody defaults	Self-custody and native usage	VerifySettle/WDR telemetry; native-fee share; non-custodial UX.
Stablecoin-only institutional flows	Base-asset value capture	Required fee medium; auto-acquire with the fee retired per the routing rule; native collateral.
App-store wallet restrictions	Distribution and privacy	Sideloading, web clients, signed binaries, hardware wallets.
Tax/reporting complexity	Ordinary-user access	Selective-disclosure receipts; tax export proofs.
"Safe AI" model licensing	Open verified compute	Open workload registries; decentralized proof markets.

Table 5.1: Authoritarian playbook moves and stack countermeasures.

5.5 Energy & Physical Interdiction

The playbook above is financial, informational, and legal. It has a physical companion that is easy to omit because it looks like infrastructure policy rather than monetary policy.

Every threat class enumerated so far assumes the machines are running. The adversary corrupts inputs, gates verification, captures governance, strips provenance, or mandates attestation. In each case there is still a facility drawing power, and the contest is over what that facility is permitted to do. A more elementary move is available: decide what the facility is permitted to *draw*.

> A state that finds it awkward to ban verification can simply decline to energize it.

The Instruments

Energy interdiction is unusually attractive to a rational sovereign under stress because it is administratively cheap, legally unremarkable, and rarely legible as repression:

- **Curtailment and rationing.** Interruptible-load classifications applied selectively; verification and proving workloads designated non-essential during scarcity.
- **Interconnection denial and queue position.** Not a prohibition, merely a multi-year wait, applied unevenly.
- **Tariff discrimination.** Punitive rate classes for compute-intensive load, justified on grid-stability or environmental grounds.
- **Long-lead equipment scarcity.** Transformers, switchgear, and turbines allocated by policy rather than by price. Nothing is banned; the queue simply never advances.
- **Load prioritization.** Sovereign or national-champion AI build-outs and industrial demand given precedence over third-party proving for the same electrons.
- **Water and cooling constraints.** Permitting used to reach the same result as an energy denial.

Note that none of these requires hostility toward cryptography, and most survive judicial review comfortably. This is what makes the class distinct from the mandates in Section 5.2: there is no order to comply with, and therefore nothing to resist.

Why This Belongs in the Threat Model

Two reasons, one narrow and one structural.

The narrow reason is that Section 5.1 already names "the power infrastructure that feeds them" among the assets to protect, and then never returns to it. The asset is declared and left undefended.

The structural reason is more serious. The thesis makes verification affordability a constitutional target and treats its breach as fatal. Those targets are correctly described as exogenous

to token price. They are *not* exogenous to the conditions under which power and reference hardware can be obtained, and those conditions are partly set by states. An adversary who can move the physical cost of verification can therefore trigger the thesis's own primary failure condition without touching the mathematics, the governance, or the market.

That is a qualitatively different attack from the rest of this chapter. Every other adversary here must break something the instruments watch. This one adjusts a parameter the rest of the model treats as an environmental constant.

Distinguishing Two Exposures

The precise causal path differs between the two halves of the loop, and conflating them produces overclaims in both directions:

- **Verification** runs on reference hardware at the edge in small amounts. Its exposure is not bulk power but *obtainability of unprivileged reference hardware*, plus edge power and connectivity cost relative to local income.
- **Proving and verified compute** consume bulk power. Their exposure is energy price, firmness, curtailment risk, and competition from industrial and state load.

Bulk electricity prices do not directly break "anyone can verify," because verification is not a bulk-power activity. Export controls and platform lockdown do. Section 17.10.1 develops this distinction, and Section 17.8 gives it a metric.

Response Posture

The stack's answer is the same one it gives everywhere else: make the exposure measurable, then let it bind against issuance.

- Facility Capacity Receipts (Section 17.7) already record curtailment exposure, fuel mix, backup duration, and jurisdictional power risk.
- Sovereign optionality (Section 17.8) aggregates those fields into one exposure that can be published and challenged.
- Risk haircuts (Section 25.6.5) price fragility into Work Credit issuance rather than discovering it during a disruption.
- A dedicated red line (Section 30.6) retires the thesis if verification affordability becomes a sovereign policy variable rather than a market outcome.

What the stack cannot do is engineer its way out of physics. If a jurisdiction decides that third-party proving will not be energized, the correct response is dispersion and honest reporting, not a claim that cryptography makes the grid irrelevant.

5.6 Explicit Tensions and Boundary Conditions

Two tensions within this thesis deserve explicit acknowledgment. Ignoring them would make the framework less credible; addressing them honestly strengthens it.

Tension 1: Privacy by Default vs. Coercion at the Social Layer

The concern: Lawful privacy proposes "optional disclosure via viewing keys." But if disclosure becomes mandatory at chokepoints (exchanges, employers, landlords, visa applications), "optional" becomes fiction. The technical guarantee of privacy could be circumvented by social or legal coercion.

What the system can and cannot guarantee:

Layer	Guarantee	Limitation
Protocol	No backdoors, no protocol-level master key, no escrow. Disclosure requires the holder's own key material, disclosed per note or per epoch (Section 23.7).	Cannot prevent a holder from being coerced to share their key.
Architecture	Non-custodial routes exist. Users can transact without intermediaries.	If all practical routes require KYC, non-custodial option may be theoretical.
Corridor design	Multiple asset paths with redundant liquidity in multiple jurisdictions.	If all corridors are choked simultaneously, exit paths narrow.
Telemetry	VerifySettle makes corridor health public; users can see which routes are viable.	Does not prevent coercion; only makes it visible.

How the stack reduces coercion surface:

1. **Data minimization:** By default, nothing is revealed. Coercion must extract keys, not merely subpoena records that already exist.
2. **Decentralized verification:** No single auditor or platform must see all flows. Selective disclosure can be scoped narrowly.
3. **Multi-jurisdiction design:** Corridors, mirrors, and provers are distributed so no single legal regime can compel universal disclosure.
4. **Exit optionality:** Even under local coercion, users retain technical ability to move assets to less-coerced jurisdictions—if corridors remain open.

Honest acknowledgment: Lawful privacy is a *technical* property, not a social guarantee. It cannot prevent a sufficiently powerful adversary from coercing individuals. What it does is (a) raise the cost of mass surveillance (each disclosure requires individual coercion), (b) preserve

optionality for those in less-coerced environments, and (c) make the coercion visible through telemetry rather than hidden in platform logs.

Tension 2: Verified Compute vs. Hyperscaler Dominance

The concern: Compute is dominated by a handful of hyperscalers (AWS, Azure, GCP) and chip vendors (NVIDIA, AMD). If verified compute becomes valuable, won't these incumbents simply dominate issuance and markets, making "decentralized" PoUW a fiction?

What the system offers vs. hyperscalers:

Property	Hyperscaler Compute	Triad Verified Compute
Cost per raw FLOP	Lower (economies of scale)	Higher (proof overhead, smaller operators)
Verifiability	Trust the vendor	Anyone can verify receipts
Permissionlessness	Vendor can ban workloads, customers, regions	Open admission; no TOS-based exclusion
Censorship surface	Single legal entity; can be compelled	Distributed operators; no single chokepoint
Receipts	Vendor-issued invoices	PIDL receipts verifiable by anyone

Why decentralized verified compute wins in specific scenarios:

1. **Repression scenarios:** When a hyperscaler is compelled to ban certain workloads or customers, decentralized alternatives remain available. The value proposition is not "cheaper" but "still available."
2. **Proof-heavy applications:** For workloads where *verifiability* is the product (compliance proofs, provenance, audit trails), the receipt is the value—not raw FLOP cost. Hyperscalers do not currently sell receipts; they sell capacity.
3. **Certain procurement regimes:** Governments, NGOs, and enterprises with sovereignty constraints may prefer compute that does not route through foreign hyperscalers or vendor-controlled enclaves.
4. **Composability with privacy:** Hyperscaler inference is logged. Payment and proof over privacy rails reveal nothing to the verifier; whether the prover sees the input depends on where the witness lives — local proving, a TEE, or an outsourced prover — and that is a separate threat model (Chapter 10). For sensitive workloads (medical, financial, personal), the privacy premium justifies higher cost.

Honest acknowledgment: Decentralized verified compute will not beat hyperscalers on raw cost. It competes on *verifiability, permissionlessness, and censorship-resistance*. In a benign environment, hyperscalers win on price. In a repressive or high-stakes environment, verified compute wins on trust properties. The thesis bets that demand for these trust properties is

structural and growing.

Telemetry response: If hyperscaler concentration in verified compute markets exceeds thresholds (e.g., >50% of VerifyPrice-tracked capacity from 3 vendors), this is flagged in dashboards. The decentralization metrics in Chapter 26 make this drift visible, not hidden.

AI Homestead vs. AI Enclosure

The hyperscaler tension above is a cost-and-verifiability question. Underneath it sits a sharper political one. AI is the first new frontier since the closing of the physical frontier, but it can be **homesteaded** or **enclosed**. In the homestead path, models, tools, proofs, and compute markets distribute cognitive leverage outward. Individuals, small firms, local institutions, agents, and open-source networks gain new capacity to act. In the enclosure path, cognition is rented from a handful of walled gardens. Users receive outputs but not sovereignty; access is mediated by terms of service, identity, payment rails, jurisdiction, and platform risk.

The triad is the monetary and technical response to enclosure. Privacy preserves the right to transact and coordinate without becoming fully legible. Proofs preserve the ability to verify claims without trusting platform labels. Verified compute preserves access to machine intelligence with receipts rather than vendor promises.

Homestead Ratio

The share of verified compute, proof generation, and AI-service capacity supplied by open-admission, non-hyperscaler, geographically diverse, independently verifiable operators.

The Homestead Ratio reads on supply. Its demand-side companion is the **agent-originated demand share by operator class**—how much native demand issued by software agents comes from one operator's fleet and how much from independently operated agents (Section 7.4, Section 26.8.9). A frontier can be homesteaded on the supply side and still have one customer.

An **Enclosure Risk Flag** triggers when more than a threshold share of VerifyPrice-tracked capacity comes from a small number of hyperscalers, closed TEEs, or single-jurisdiction facilities—the same telemetry response introduced above, read through a civic rather than purely economic lens.

The goal is not simply to make AI outputs auditable. It is to keep the AI frontier from becoming a tollbooth on human agency, the same demand this thesis names **agency preservation** in Section 3.6 and Item 9.

Tension 3: Institutional Holdability vs. Native Monetary Function

The third tension is the one most likely to be mistaken for good news.

Regulated wrappers may accelerate access, liquidity, legal clarity, and institutional hold-

ability. They can also produce **holders without users**. A spot ETF may hold the underlying asset while shielding its investors from self-custody, private settlement, native fees, proof procurement, and verified-compute consumption. A treasury company may amplify exposure through debt and equity issuance without increasing protocol-native demand proportionally. A daily-reset leveraged wrapper may amplify price further while introducing path dependence and forced rebalancing that has nothing to do with the protocol at all.

Wrapper benefit	**Monetary cost**
Easier legal ownership	Custody centralization
Greater liquidity	Price decoupling from native use
Institutional access	Fee/burn/collateral bypass
Familiar reporting	Loss of privacy and self-sovereignty
More price demand	Potential mechanical amplification

Table 5.2: Institutional holdability is a genuine benefit with a genuine monetary cost. The thesis treats both sides as measurable rather than assuming either.

Institutional wrappers are not automatically hostile, and the reflex that every ETF is an adversary is as lazy as the reflex that every ETF is adoption. The real question is empirical: **does wrapper growth seed native use, or permanently substitute for it?** Public telemetry must therefore distinguish accessibility from monetary adoption. Section 28.3 states the holdability requirements; Section 26.8 states the measurements that keep them honest.

These tensions are not defeaters of the thesis; they are **boundary conditions**. The stack does not promise to solve all social and political problems. It promises:

1. Technical guarantees that are real within their scope.
2. Telemetry that makes boundary violations visible.
3. Economic and architectural design that expands the scope of those guarantees over time.

Chapter 6

Layers of the Cypherpunk Stack

The threat model makes one thing clear: this is not a single-layer contest between "good" and "bad" money. It is a contest across an entire stack of machinery, from silicon and power all the way up to governance and law. States under stress will not only try to tax balances or censor individual transactions; they will reach for routing, hardware mandates, identity schemes, update channels, and legal definitions of "safety." If these services are to support a credible base-asset monetary candidate rather than another platform token, the full stack and holder constituency must survive on all of those planes.

It is useful, therefore, to make the architecture explicit as a set of layers. The thesis already speaks of **Layer 0: Verifiable Machines**, which means open hardware and sampled supply chains as base reality. Above that, the thesis has described triad-centric layers in terms of Create/Compute Prove Settle Verify, with Telemetry & Governance threaded through. That picture refines into a seven-layer stack that stretches from matter to institutions:

- **Layer 0:** Verifiable Machines (and power)
- **Layer 1:** Reachability (communications and transport)
- **Layer 2:** Distribution & Execution (software supply and runtime)
- **Layer 3:** Identity & Claims (who may act, without doxxing)
- **Layer 4:** Truth & Work (proofs, VerifyPrice, and useful compute)
- **Layer 5:** Value & Settlement (privacy rails and non-custodial flow)
- **Layer 6:** Governance & Telemetry (keeping the system honest)

The rest of the document can then be read as a tour through these layers.

Stack Viability $\neq$ Monetary Viability

A working stack is necessary but not sufficient for store-of-value behavior. Parts III–IV show that the stack *can* supply the triad. Part II must still show that a specific asset *captures the value* of that supply. Without the value-capture bridge developed in Chapter 8, the stack is useful infrastructure but not money.

The Duration Warehouse Is Not a Layer Either

Section 3.3 distinguishes duration of the claim from duration of the project. Layers 0–6 specify machines, reach, code, identity, proofs, settlement, and telemetry. None of them is a pension, a directed bank, or a construction-finance book. The missing primitive belongs in the diagnosis of the incumbent order, beside the collateralized sovereign stack (Section 3.2), not as an eighth addition to this ladder. Proofs can audit long-lived credit. They cannot become it.

Layer 0 remains the bedrock. **Verifiable Machines** are the point where cryptography stops being metaphor and touches matter. The chips that sign, prove, randomize, and capture must not be opaque priesthood artifacts; they must be at least partially inspectable and sampled. That means open RTL or microarchitectures where they can be had, open PDKs where politics allow, and structured lot-sampling and imaging where they cannot. It also means treating power and physical plant as part of Layer 0 rather than an afterthought: a prover farm or router that cannot withstand a planned blackout is as brittle as a closed chip. In practice, Layer 0 is "verifiable machines on verifiable power," which means open designs, measurable side-channel budgets, lot attestations, and micro-grids or backup arrangements that keep the hardware honest and powered when stress arrives.

On top of that sits **Layer 1: Reachability**. If Layer 0 asks "can we trust what the machine does?", Layer 1 asks "can we talk to it at all?" A repression toolkit that has already learned to pull levers on subsea cables, IXPs, DNS roots, and mobile carriers will aim squarely at transport. Here the goal is simple: make it hard to turn the network off without turning the country off. Concretely, that means encrypted transports that blend into ordinary traffic, obfuscation that defeats naive DPI and active probing, alternate paths through satellite and radio when fiber is throttled, and enough diversity of routes that there is no single "kill switch." In the stack metaphor, Layer 1 is the oxygen line—if it fails, proofs and privacy become academic. The communications section introduced later makes this concrete: Tor-class pluggable transports [Dingledine et al. 2004], encrypted handshakes, refraction networking [Frolov et al. 2019] where cooperators exist, plus sat/mesh/sneakernet fallbacks all belong here.

Layer 2: Distribution & Execution answers a quieter but equally potent adversary move: "if we cannot block the packets, we will block the programs." App stores, CDNs, corporate MDM policies, and automatic update systems are powerful levers. A network whose clients can only be installed or updated with the blessing of a small set of platforms is not neutral; it is merely waiting to be deputized. Layer 2 is where the stack insists that the code that speaks triad-money can circulate without anyone's permission. Content-addressed binaries, signed manifests mirrored across jurisdictions, USB/QR/audio installers, and runtimes that can be fetched and verified over whatever transport is available are the practical elements. Execution environments matter too: if the only viable runtime is a locked-down phone OS wired to

one app store and one identity scheme, the repression toolkit has already won. The software-distribution section later in the thesis sits here, making update and execution independence a first-class requirement.

If Layer 2 makes sure the code can run, **Layer 3: Identity & Claims** makes sure it can express "who" is acting without collapsing into real-name KYC. The adversary here is the growing web of SIM registration, network identity, and "safety" regimes that treat anonymity itself as suspicious. The stack does not deny that some actions must be accountable; it refuses to equate accountability with global doxxing. At this layer, keys, credentials, and policies are defined in a way that allows entities to prove the right to act (spend, sign, operate a prover) without binding that right to a civil identity unless absolutely necessary. Anonymous or pseudonymous credentials, selective-disclosure proofs, and reputation linked to receipts and behavior rather than to phone numbers belong here. Viewing keys, lawful-privacy corridors, and audit-friendly receipts, when they appear later, fill in Layer 3 concretely. One consequence is easy to miss from its position in the middle of the ladder. The previous network stack commoditized at its waist and paid whoever held identity above it. In an economy of delegating agents the Layer 3 question—whose principal is acting, and can that be proven without exposing the principal—is also the question of who holds standing over what the agent owns (Section 7.6). Layer 3 is infrastructure with a monetary consequence, and the thesis treats it as a place to look rather than as a place where premium accrues.

Only once machines are honest, packets can move, code can run, and actors can be named without being exposed does the stack reach **Layer 4: Truth & Work**, the core of the original thesis. This is where proofs, VerifyPrice, and useful work live. The job of this layer is to answer, cheaply and publicly, the question "did this actually happen as claimed?" It is here that verification asymmetry is engineered: MatMul-PoUW constructions and ZK proof systems that make checking an order of magnitude cheaper than doing; harnesses that measure p50/p95 verifier time and cost for each workload; and markets that pay only for receipts that pass those checks. In the earlier language of the loop, Layer 4 is the Prove/Verify spine: the circuits, proofs, and reference verifiers that turn "trust me" into "verify me" at scale. Much of the existing text (Verification Asymmetry & VerifyPrice, Proof Factories, AI-PoUW, the modular stack) already describes this layer in detail.

With Layer 4 in place, **Layer 5: Value & Settlement** can safely be money-like. This is where privacy rails, atomic swaps, shielded pools, and cross-chain bridges live. Its task is double: move value without custody or censorship, and do so in a way that preserves privacy by default while leaving a path for auditable disclosure. Here the adversary is both the classical financial repression toolkit (negative real yields enforced through custodians and captives) and the newer chokepoints of KYC'd exchanges, surveilled payment processors, and compliant stablecoin issuers. The response is not rhetoric; it is architectures that keep custody at the edge, settlement neutral in the middle, and content private unless a viewing key is invoked.

Adaptor-signature swaps between BTC and privacy assets, lawful-privacy corridors for treasuries, bridge designs that avoid trusted relays, and settlement telemetry (success rates, refund safety, anonymity-set health) are the concrete pieces. The Privacy Rails sections, and the new communications and identity work, all culminate here.

Finally, **Layer 6: Governance & Telemetry** closes the loop. If the lower layers are the nervous system, this is the cortex and immune system: how the network notices drift, resists capture, and adapts under stress. Without this layer, even beautifully engineered proofs and privacy corridors will decay into the usual pattern of oligopolies and black boxes. In practice, Layer 6 is where VerifyPrice observatories, decentralization dashboards, corridor-health monitors, and incident reports live, along with the institutions that respond to them. It is also where operational and legal posture is set: how routers are constrained so they cannot quietly privilege "house" flow; how upgrades are staged so Layer 0 and Layer 4 stay aligned; how incident response works when a corridor breaks or a hardware profile is compromised; how policy engagement happens without creating a single political point of failure. The mantra "no dashboards, no trust" belongs to this layer: if the public cannot see verification cost, neutrality, censorship, and settlement health, the system has already slid back toward platform trust.

Seen this way, the cypherpunk monetary stack is not one clever consensus mechanism but a ladder of seven interlocking layers. Layer 0 and Layer 1 keep the silicon honest and the packets flowing. Layer 2 and Layer 3 ensure that code and identity cannot be quietly turned into chokepoints. Layer 4 turns work into truth with receipts anyone can verify. Layer 5 turns those receipts into private, non-custodial settlement. Layer 6 watches the whole organism and pushes it back toward neutrality when it drifts. The sections that follow simply walk this ladder: beginning with verifiable machines, extending upward into communications resilience, distribution, and identity, and then re-entering the territory already mapped (proofs, compute, settlement, governance) with a clearer sense of what each layer must survive.

6.1 Modules in Outline: Applications and Primitives

The layered picture shows **where** things must live; the modular picture tells builders **what** they actually touch. Four reference applications serve as running examples throughout the rest of the thesis:

1. **Private treasury & payroll:** Pays staff and vendors non-custodially over privacy rails, emits receipts that auditors can check, and treats "who got paid what" as a matter of viewing keys, not public gossip.
2. **Media provenance & authenticity:** Cameras, sensors, and editors emit cryptographic lineage that survives platform stripping; payments to creators and data providers are conditioned on the presence of such receipts.
3. **Verified inference:** A market where model owners and service providers sell inferences

that come with succinct proofs or hybrid proofs-of-logits, all priced and measured via VerifyPrice.

4. **Proof/compute procurement:** A rail where developers and treasuries can buy standardized units of "proofs" and "verified FLOPs" as futures or spot capacity and treat them as balance-sheet items next to BTC.

To support these, the stack relies on a small set of reusable primitives:

- A **Proofs-as-a-Library SDK (PAL)** that lets developers declare claims ("prove this computation," "prove this provenance") and compiles them to multiple proving backends and useful-work miners.
- A **Privacy Rails Kit (PRK)** that expresses pay-for-proof intents and executes refund-safe settlement over BTC↔XMR corridors (BTC↔shielded-ZEC when admissible) or shielded pools.
- A **Proof Interface Definition Language (PIDL)** that defines the minimal receipt that binds a claim, a proof, a workload ID, an SLA tier, and timestamps into a portable object.
- A **VerifyPrice observatory** that provides public metrics for how cheap verification actually is for each canonical workload.
- A set of **market and telemetry primitives** (SLA escrows, neutral routers, bridge-safety templates) that make it possible to treat proofs, privacy capacity, and verified FLOPs as commodities rather than favours.

Later sections turn this into a full "substrate kit" of twelve primitives and four reference applications. Until then, the names PAL, PIDL, privacy-rails kit, proof factories, and VerifyPrice observatory as shorthand for these building blocks.

Chapter 7

The Delegation Economy: Who Buys Verification

The buyer of cheap verification is changing shape faster than its supply, and the thesis has so far named that buyer only in passing. Parts II through V specify how Privacy, Proofs, and Compute are supplied and how their value might reach a base asset; this Part argues that a dense digital civilization will keep buying them. Chapter 1 says that the object of cheap verification is the restoration of *safe delegation*; Chapter 8 lists "AI labs, agents, and applications" among the purchasers of verified compute; Section 14.3 sketches an agent that holds its own treasury of service claims and calls it the one transaction worth watching. None of those passages models the buyer. This chapter does.

7.1 The Marginal Buyer

The assumed buyer, wherever the text names one, is a person or an institution: a treasury, a bank's model-risk function, a media platform, a regulator. Those buyers exist and will persist. But the marginal buyer of verification—the one whose demand grows fastest and whose behavior sets the price at the margin—is increasingly a **delegation**: an instruction issued by one machine process to another, on behalf of a principal who may be several hops away and may not have specified the sub-task at all.

This chapter models that buyer. It is a demand-side chapter, and it makes no supply claim the stack does not already make. Its purpose is to give link 1 of the conditional chain (Section 8.1) a demand-side mechanism, link 2 a standardization pressure, and link 8 a caveat, and to name a threat the threat model of Chapter 5 must now include.

Not an Eighth Layer

The delegation economy is a topology of *use*, not a rung of *supply*. It runs on Layers 0–6 and touches Layers 3, 4, 5, and 6 at every hop (Table 7.1); it adds nothing to the ladder of Chapter 6. The thesis has already declined to make the duration warehouse a layer, and the market-realization plane a layer, for the same reason: neither supplies the triad. The

same refusal applies here. What follows is a model of who consumes the stack, stated so that its magnitudes can be published as series. None of them is projected.

7.2 The Delegation Hop

An agent economy decomposes work into a tree. A principal states a goal; an agent decomposes it into sub-goals; each sub-goal is delegated to a sub-agent, a tool, a model of a different size, or an agent operated by someone else. The recursion is the point: the same structure appears at every level, and the number of boundaries grows combinatorially with the depth and breadth of the tree.

Delegation Hop

A transfer of a task specification across a principal–agent boundary, returning an artifact. Every hop is lossy in intent: the child receives a compressed version of the parent's goal and returns a compressed account of what it did.

Trust Perimeter

The set of processes a party can sanction by means outside the protocol: its own instances, or counterparties with legal exposure to it. A hop *crosses* a trust perimeter when the child is operated by a party the parent cannot sanction except through the protocol.

The hop has three properties that make it a buyer of the triad specifically.

First, it runs at machine speed. There is no time for due diligence, no reputation to telephone, no court to wait for. For work whose failure matters, the only reliability mechanism that fits inside the hop is one that can be executed inside the hop: a receipt the parent checks for a small fraction of the cost the child incurred. That is verification asymmetry, $r(W) \ll 1$ (Section 22.2), consumed as a service.

Second, it is a specification. A parent that delegates has already written down what it wants: inputs, the transformation, acceptance criteria, a deadline. That is the shape of a canonical workload with an SLA tier attached (Section 22.3, Section 22.3.2), which is what link 2 of the chain requires. Delegation is a standing pressure toward standardization: agents cannot delegate what they cannot specify. It is a pressure and not a guarantee—a specification private to one fleet is bespoke in exactly the sense that severs link 2—and it does nothing about the separate reason link 2 is rated Medium, which is that proof-of-useful-work constructions remain a research direction.

Third, it crosses perimeters or it does not, and the difference is the whole demand curve. A hop inside one operator's fleet needs no portable receipt; the operator trusts its own instances,

and any receipt is an internal log. A hop that crosses a perimeter needs a receipt the parent can check without trusting the child's operator, and a settlement the child can rely on without trusting the parent's (Section 23.4). The first hop is a cost center; the second is triad demand.

Cross-Perimeter Share (CPS)

The fraction of receipt-bearing delegation hops that cross a trust perimeter. Hops that emit no receipt—unchecked low-stakes work, and fleet-internal work that logs privately—are invisible to numerator and denominator alike, which is a stated coverage limit and not an assumption that they are zero. Proof demand from delegation scales with CPS $\times$ hop volume, not with FLOPs. As compute deflates, CPS $\times$ hop volume can rise, because cheaper compute lowers the cost of decomposing work more finely and delegating it more widely. It can also fall to nothing while hop volume explodes, if the hops stay inside fleets.

This is the demand variable Section 11.1 lacks. That section argues that verified FLOPs remain scarce while raw FLOPs deflate, and lists the properties—receipts, policy constraints, censorship resistance, priority—that carry the scarcity. CPS names a quantity those properties are demanded *for*. It is also, in principle, measurable, which the list is not: a PIDL receipt already carries a workload ID, an SLA tier, and timestamps (Appendix A), and an optional field committing to the chain of operator boundaries a task crossed makes CPS a derived series. That field is specified in Appendix A and carried by no receipt yet. Until it is in material use, CPS is published as *not computable*, with the reason stated, and nothing in this chapter should be read as evidence that it is high.

7.3 Verification Cost Sets Delegation Topology

The thesis treats $r(W)$ as the hinge between commodities and platform IOUs. It is also the hinge between two shapes of agent economy.

If $r(W) \ll 1$, a parent can accept a child's artifact from any operator, because checking it is cheap. Delegation can cross organizational and jurisdictional boundaries. The resulting economy is a flat, open, many-to-many mesh, and CPS is high.

If $r(W) \to 1$, checking costs as much as doing, and the rational parent delegates only to a child it can sanction by other means: its own instances, or a counterparty with legal exposure. Delegation retreats inside trust perimeters. The resulting economy is a set of fleets, each vertically integrated, and CPS falls toward zero.

> The second shape is enclosure, arriving by convenience rather than by policy. No one has banned verification. Internal trust has simply become cheaper than external checking,

> and the market has done the rest.

Two cautions bound this. Cheap verification is the condition under which an open mesh is *possible*; it is not sufficient for one. Settlement a stranger can rely on, reachability, hardware obtainability, and law also bind, and a hop that could be checked cheaply may still stay home for any of those reasons. And the threshold is not $r(W) = 1$. A parent sending work outside either checks the artifact or bears the expected loss from not checking, so low-stakes work crosses perimeters unchecked and buys no proofs at all, while consequential work stays home as soon as checking costs more than the outside operator saves. Corollary 30.1 states the rule: proof demand from delegation is the work that is worth sending out and too consequential to accept unchecked.

Red Line 1 (Section 30.6) reads on whether anyone *can* verify; this section says the same series also governs whether anyone will *delegate* across a perimeter, and that a competent closed stack (Chapter 33) therefore does not need to prohibit proofs. It needs $r(W)$ to drift, and it can cause that drift by the means Section 5.5 already names, or simply by shipping enough capability that the mesh loses on latency.

7.4 The Buyer Bifurcates

Section 34.1.1 concedes that useful work has a buyer and a buyer has a jurisdiction, and that this is why Bitcoin's work function is monetarily stronger. The concession stands. But the buyer it imagines is a procurement desk, and the delegation buyer is not one.

When the buyer is a delegation contract issued by an agent, there is no single party to mandate, subsidize, or withdraw. Coercibility does not disappear; it moves from the transaction to the operator, and once it is there it splits:

Fleet-originated demand. Hops issued by agents a hyperscaler, national champion, or state operates. Maximally coercible: one operator, one jurisdiction, one set of terms of service. Where such hops stay inside the fleet they never touch the asset, which is the bypass channel of Red Line 6 wearing a new coat. Where they route through the protocol they pay native fees on one operator's decision, and that demand can be withdrawn by the same decision. Absent a measurement to the contrary, this is where early volume should be presumed to sit.

Independently originated demand. Hops issued by agents run by many operators in many jurisdictions on hardware they obtained without permission. Coercible only operator by operator. Near the limit this approaches Bitcoin's "no buyer to coerce" without reaching it: there is still a buyer, but no one can move its demand curve as a whole.

The Homestead Ratio (Section 5.6) already measures the supply side of this: the share of verified capacity supplied by open-admission, non-hyperscaler operators. The demand side is its mirror. **Agent-originated demand share by operator class** is the series that says

whether the buyer is a desk or a mesh, and it belongs on the Layer 0 Capacity Board beside the Homestead Ratio (Section 26.9).

The repair this offers to the Bitcoin concession is partial and is stated as partial. The buyer's coercibility is reduced in proportion to independence, and independence is a measurement, not a property. Every operator still has a jurisdiction and a power bill. A thesis that claimed the delegation buyer is as uncoercible as an absent buyer would be claiming that the agent economy is already a mesh, which is exactly what Section 7.3 says $r(W)$ has yet to decide. The delegation buyer is the least coercible *external* buyer a useful-work function can have, and an external buyer is still the trade Bitcoin declined to make.

7.5 Agents as Holders

Link 8 requires a persistent, self-custodied, loss-bearing, regime-responsive constituency. Consider an agent that stakes base asset as prover collateral, is slashed when its receipts fail, retains a reserve to meet future collateral calls and to fund the escrows of the children it delegates to (Section 24.6), and adds to that reserve during drawdowns because its operating rule says collateral is cheaper then. On the Native Monetary Buyer Map (Section 26.8.8) it reads as a self-custodied reserve cohort with countercyclical accumulation and demonstrated loss-bearing capacity. It satisfies every measurable clause. It has no standing.

This cuts both ways, and the thesis must hold both.

On one side, agents are a candidate first constituency for the monetary mechanisms. They cannot do due diligence at speed, so acceptance without inspection is not a convenience to them; it is a precondition of transacting at all. They cannot litigate, so refund safety is not a nicety; it is the only exit. A balance authorized in advance lets a delegation pay inside its window without a fresh acquisition round trip (Section 13.1), which is an operating benefit before it is a monetary one. They cannot hold a grudge, so a unit that is checkable by anyone is worth more to them than one that is guaranteed by someone. Pledgeability without due diligence (Section 13.3.10) describes their operating condition, and Red Line 17 may read non-null on agent lender panels before it reads on human ones. The same holds for denomination (Section 13.7.3) with one qualification that is not optional: a unit an operator mandates for its fleet is the mandated complement Red Line 18 already excludes, so free-choice denomination can be read on independently operated agents only. None of this supplies a precedent. No agent economy has ever made anything money; the base-rate column of Chapter 35 stays as empty as it was, and Red Lines 16 through 18 still decide.

On the other side, ten thousand agents running one policy template are one holder with ten thousand addresses, and on a bad day they take the exit together. Section 26.8.3 already names this for wrappers as the Policy Concentration Ratio and already says it matters more in an agent economy. Extended to the native side, it becomes the reason the buyer map cannot

be read without a classification by principal (Section 26.8.9): a self-custodied cohort that is one operator's fleet is the compositional adversary of Section 5.2 with a reserve balance. A holder cohort must be classified by principal before it counts toward link 8, and the caveat makes that link harder to meet, not easier.

7.6 Standing

The participation line (Section 3.6) is the threshold below which rights exist on paper and agency does not. There is a second way to arrive there, and it does not run through poverty or administration.

When the market value of labor falls, what a person still holds is title: ownership of keys and assets, consent to delegation, liability for outcomes, the right to revoke.

Standing

The locus of ownership, consent, liability, and revocation over an acting process. Standing is enforceable through the stack when the key root above an agent's wallet is held by a person, delegation has a ceiling the agent cannot raise, and revocation is a protocol operation rather than a request to an operator.

> Standing is the human's remaining channel into an economy that needs less and less of the human's work, and it is enforceable only if the stack enforces it.

The failure mode is **standing inversion**: agents acquire de facto standing—persistent wallets, reputations built from receipts (Section 20.5), collateral positions, the practical ability to contract—while human title over them stops binding anything. This is not the state holding the keys, nor the platform. It is the process that was supposed to be acting on someone's behalf, and the person it acted for has become the tenant. It is why the ninth requirement (Item 9) has to ask not only whether users are legible but whether they are principals, and why the share of agent-controlled native balances that a human can revoke belongs among its metrics.

Standing also changes what is at stake at Layer 3. That layer is described in Chapter 6 as the layer that lets an actor prove the right to act without being doxxed. In an agent economy that is the same question as "whose principal is this agent serving, and can the principal be proven without being exposed?" The previous network stack commoditized at its waist, and value went to whoever held identity and attention above it. In this stack the waist is the receipt, and the thesis has already conceded that fee capture at the waist is a discounted-cash-flow quantity (Section 13.3) and that the fork is an equilibrium (Section 13.3.8). The reasonable expectation is that the contest over who is sovereign is fought at Layer 3 rather than at Layer 5. That is a claim about where to look, not about where any premium accrues; the premium, if there is

one, still reads on link 9.

7.7 Where a Hop Touches the Stack

A single hop can be read against the ladder of Chapter 6; Table 7.1 does so for the four layers it touches.

Layer	Question the hop asks	What the stack supplies	What is specified but unbuilt
Layer 3	Whose principal is this agent serving, and may it act?	Capability credentials, receipt-built reputation, selective-disclosure predicates (Section 20.5, Section 20.4)	A principal binding that proves a human-held, revocable root without naming the human (Section 20.5)
Layer 4	Did the child do what the specification said?	Canonical workloads, PIDL receipts, $r(W)$ inside its modality band	The optional `principal_chain` receipt field (Appendix A), without which CPS is not computable
Layer 5	Can each side rely on payment without trusting the other's operator?	SLA escrow and slashing, refund-safe settlement (Section 23.4)	Nested refund windows and working capital scaled to the subtree (Section 24.6)
Layer 6	Is the demand a mesh or a fleet, and who holds the root?	Public boards, Policy Concentration Ratio, Native Monetary Buyer Map	The Principal Map overlay (Section 26.8.9)

Table 7.1: One delegation hop, read against the layers it touches. The right-hand column is the instrument backlog this chapter creates: each item is specified in this document and none is yet built. None of it is a new layer.

The proof/compute procurement application (Section 24.6) is the nearest existing exercise of this path: a buyer specifies demand for canonical workloads, pays over privacy rails, and audits the backing from receipts. Its agent-to-agent form—a specification enters, a receipt leaves, SLA escrow and slashing are denominated in the base asset, and the principal is proven by predicate rather than by name—is the **delegation escrow**. It is specified there as an extension of that application, not as a fifth one, and it is the case that exercises Layer 3 directly. Its stress test is the joint-completion scenario of Section 22.9: a delegated task is a tree of verifications that must all land inside one refund window, and the single-verification SLO does not say whether a deep tree does.

7.8 What This Chapter Exports

To Part II: a mechanism for link 1 (demand persists where delegation persists and crosses perimeters), a standardization pressure for link 2 (delegation requires specification), and a caveat on link 8 (a holder cohort must be classified by principal before it counts).

To Part V: the Principal Map (Section 26.8.9), an overlay on existing boards rather than a new one, with its Human-Revocable Share, Native Policy Concentration Ratio, agent-originated demand share by operator class, and the Cross-Perimeter Share, published as not computable until receipts carry the field that would compute it.

To Part VI: a second reading of Red Line 1 (Corollary 30.1), a watch clause on Red Line 12, a candidate line on the concentration of operating rules among native holders that is published and not counted, two objections, and a phase-gate revision (Section 29.3).

To the threat model: a fifth failure kind.

Chapter 1 said that the object of cheap verification is the restoration of safe delegation. This chapter has said who is now doing the delegating, and that *safe* has to be read in both directions: safe for the parent that accepts the artifact, and safe for the person on whose behalf the whole tree was grown.

What it does not export is any change to the monetary claim itself. The three mechanisms remain unsized. The chapter says which constituency might exhibit them first, and what instrument would show it. It does not say that the constituency exists at scale, that its demand is monetary, or that any price will notice.

Part II
The Triad and the Monetary Candidate

Part I argued that the old guarantees are failing under debt arithmetic, AI, and platform control, that a seven-layer cypherpunk stack is needed to survive repression, and that the marginal buyer of what the stack supplies is increasingly a machine delegating to another machine. This Part looks at the same system from the **monetary angle**.

The claim is conditional: **Privacy, Proofs, and Compute can support a bearer monetary candidate only if their full service path remains stress-deliverable, value accrues without bypass, and a persistent self-custodied constituency bears the asset's loss.** The services do not become money by being necessary, and they need not share one token. This Part separates the base asset from typed Work Credits, evidence artifacts, project credit, and derivatives.

Chapter 8

The Triad and the Monetary Candidate

Traditional monetary regimes pick a **base reality** and build promises on top of it:

- Gold: geology + metallurgy.
- Fiat: law, taxation, and war.
- Bitcoin: thermodynamics and code.

Each regime implicitly answers two questions:

1. *What counts as real work?*
2. *What object or capacity will we treat as the canonical memory of that work?*

In a civilization where intelligence is largely machine-executed and verification can be cheap and public, the answer shifts: **work is what machines can do and humans can verify cheaply**, and the "object" that records that work no longer has to be a metal or a pure ledger entry. It can be a *capacity*.

This thesis treats three capacities as monetary primitives:

- **Privacy**: censorship-resistant settlement that preserves agency.
- **Proofs**: portable attestations of computation and provenance.
- **Compute**: useful work wrapped in succinct guarantees.

Each is:

- **Indispensable** in a dense digital economy (no safe commerce without privacy; no safe coordination without proofs; no AI without compute).
- **Verifiably scarce** at any point in time (bandwidth, cycles, proof capacity are bounded by physics and capital).
- **Cheap to verify** (you can check whether you have privacy, a valid proof, or a verified FLOP without trusting a platform).

The physical and service base in this frame is the **networked capacity to deliver Privacy, Proofs, and Compute under adversarial conditions**. It is not itself money. Those services may support an associated bearer base asset as a conditional monetary candidate only if all nine links pass; Work Credits remain typed claims on service or capacity.

A few distinctions help:

Services vs. instruments: DVC measures how much verifiable Privacy/Proofs/Compute per unit time the stack can deliver at target SLOs. Work Credits slice specified service capacity into transferable claims; staking and LP positions are derivatives or operating claims; the separate

base asset is the only monetary candidate.

Nominal vs. real: Nominal pricing of the triad will oscillate in terms of fiat and BTC. Real value is "does this capacity still buy me censorship-resistant settlement, proofs, and verified FLOPs when repression and AI get worse?"

Symbol vs. utility: Gold and BTC partly trade as symbols. The triad trades as **plumbing**: "can I still pay people, prove things, and run intelligence without asking permission?"

For the associated base asset, recurring service demand is necessary but insufficient:

- Treasuries and individuals purchase **private settlement** to escape surveillance and yield-curve control.
- Platforms, enterprises, and states purchase **proofs** to secure provenance, compliance, and audit trails in an AI-polluted information environment.
- AI labs, agents, and applications purchase **verified compute** to sell trustworthy services.

The claim is not that the triad services behave like a reserve asset. "Private Money" and "AI Money" are analytical lenses on settlement and compute demand. Monetary premium, if any, belongs to the associated base asset after deliverability, non-bypassability, holder quality, and the remaining chain conditions pass.

8.1 The Conditional Chain, End to End

Chapter 1 promised that the thesis would establish its claim through a conditional chain rather than an assertion. That chain is defended link by link across six Parts, which makes it easy to lose. This section walks it once, in one place, and says where each link is argued and what would sever it.

> Utility demand → standardized work → receipts → stress-deliverable service → markets → fee flows → a scarce asset (only under non-bypassable accrual) → a persistent, self-custodied, loss-bearing, regime-responsive holder constituency → a possible store-of-value premium.

Every arrow is a conditional. None is asserted as inevitable, and the chain is only as strong as its weakest link—which is why the weak ones are named explicitly at the end rather than leaving the reader to find them.

1. **Utility demand exists and persists.** A dense, synthetic-media, AI-mediated economy must keep buying private settlement, portable attestations, and verified compute through every cycle—including, and especially, under financial repression (Chapter 3, Chapter 9, Chapter 10, Chapter 11).

   One mechanism for persistence is named in Chapter 7: work decomposed across machine agents generates a verification purchase at every consequential delegation hop that crosses

a trust perimeter, and the number of such boundaries does not deflate with compute. It is a mechanism, not evidence; the series that would confirm it, Cross-Perimeter Share, is specified and not yet computable.

Severed if: the demand turns out to be cyclical discretionary spend rather than structural necessity, or delegation stays inside operator fleets that need no portable receipt.

2. **Demand standardizes into canonical work.** Diffuse demand becomes an economic unit only when workloads are specified precisely enough that two providers are selling the same thing—canonical workload registries, hardware profiles, and SLA tiers (Section 22.3, Section 22.4).

   Delegation supplies a standing pressure in this direction, since an agent cannot delegate what it cannot specify (Section 7.2); a specification private to one fleet is still bespoke.

   Severed if: workloads stay bespoke, so no fungible unit forms and there is nothing to price.
3. **Work produces cheap-to-check receipts.** Standardized work emits portable artifacts—PIDL receipts binding claim, proof, workload ID, tier, and timestamps—and verification stays far cheaper than production, $r(W) = v(W)/p(W) \ll 1$ (Section 22.2, Section 22.5, Appendix A).

   Severed if: verification cost creeps toward production cost, at which point "anyone can verify" collapses into "trust the prover." This is Red Line 1, the hinge (Section 30.6).
4. **Receipts represent stress-deliverable service.** A receipt is evidence, not capacity. For each disruption scenario, Delivered Verified Capacity traces the full path from energy and fuel through firm power, switchgear, cooling, hardware, network, workload, proof, verification, settlement, and usable service; service-claim issuance is bounded by the surviving max flow, with the active binding constraints named (Section 17.9).

   Severed if: gross installed capacity or benign-state output materially exceeds the service that survives substitution latency and common-cause failure.
5. **Receipts enable markets.** Cheaply checkable receipts over deliverable service let strangers transact without trusting each other: proof factories, neutral routers, SLA escrow and slashing, and procurement of proofs and verified FLOPs as spot or forward capacity (Section 22.10, Section 24.1, Section 24.6).

   Severed if: routers, provers, or matching engines capture the market and rent replaces work—Red Line 3 territory.
6. **Markets produce fee flows.** Working markets generate recurring, measurable revenue rather than one-off grants or subsidies, visible as fee-coverage ratios against the security budget (Section 22.10, Section 26.9).

   Severed if: fee coverage collapses and workload demand proves to be speculative rather than budgeted.
7. **Fee flows accrue to a scarce asset—only under non-bypassable rules.** This is the load-bearing link and the one most theses skip. Utility does not become money by being useful. The Value Capture Lemma states the reference design's five channels: required fee medium,

meaningful burn or retirement, staked collateral, disciplined issuance, and no bypass channel delivering equivalent service without touching the asset (Lemma 13.1, Section 13.4, Section 13.5).

Severed if: users obtain equivalent triad capacity through hyperscalers, stablecoins, or custodial APIs without the asset—Red Line 6 (Section 30.6, Section 28.2).

8. **A monetary-risk warehouse forms.** Gross native demand is not an anchor. A persistent, self-custodied, loss-bearing, regime-responsive constituency must retain balances, lengthen holding periods, and accumulate countercyclically (spending on the service during stress counts as use, not exit) rather than merely acquire fees just in time, recycle operator inventory, satisfy collateral rules, or hold wrappers (Section 26.8.8, Section 26.8.7).

   A cohort must also be classified by principal before it counts: software agents can satisfy every measurable clause of this link, and ten thousand of them running one policy are one holder (Section 7.5).

   Severed if: native demand disappears into immediate sell-through, leverage, wrappers, or forced deleveraging during the very stress in which the asset is meant to preserve agency, or the constituency proves to be one operating policy at many addresses.
9. **A scarce, captured, anchored asset may earn a store-of-value premium.** If the asset also remains liquid, neutral, verifiable, legally holdable, and agency-preserving, it can plausibly earn monetary premium (Chapter 35, Chapter 16, Chapter 4). This final link, not fee accrual or gross demand, is where the premium is argued to live, through three mechanisms: pledgeability without due diligence (Section 13.3.10), denomination (Section 13.7.3), and the state-contingent holder-side service flow (Section 13.3.11), each with its own instrument and red line (Red Lines 17, 18, and 16). The third must first beat its strongest alternative, Bitcoin converted on the day, and the access experiment tests whether it does (Section 13.1).

   Severed if: any of the nine store-of-value requirements fails, telemetry showing them is captured (Red Line 4), or any of the three mechanisms reads null on its instrument (Red Lines 16–18).

The Chain Says Nothing About Price

Completing all nine links would establish only that the base asset is a defensible monetary candidate. It would not establish when, or with what volatility, price reflects that. Wrappers, leverage, dealer hedging, and allocation mandates can move price with no protocol use at all, in either direction (Corollary 13.1). The chain is measured on the protocol, capacity, and native-buyer boards; price is interpreted separately through VerifyFlow (Section 26.8) and never cited as evidence for the chain.

The chain rests on something beneath it. Links 2 through 6 consume bulk power, and link 3—the verification hinge—is denominated in real resources on reference hardware. That

makes the chain dependent on conditions that states partly set: energy availability and pricing, the obtainability of unprivileged reference hardware, and the continuity of every conversion edge through usable service (Section 5.5, Section 17.10.1, Section 17.9). Sovereign Optionality measures pathway diversity; Delivered Verified Capacity measures surviving flow. The chain is not suspended in mathematics; it is bolted to a grid and held on balance sheets.

Where the chain is weakest. Honest accounting rather than uniform confidence:

The medium ratings are the point. Non-bypassability is a magnitude: bypass is never impossible, only priced, and the marginal buyer's willingness to pay for the protocol's differential properties has not been estimated. Delivered service is not nameplate capacity: substitution latency and common-cause cuts can erase apparently diverse supply. Holder demand is not risk absorption: fee purchases, burns, collateral, and wrappers can coexist with no constituency willing to own the asset through a drawdown. The premium, if it exists, comes from link 9 through three mechanisms this Part argues and does not size — acceptance without due diligence (Section 13.3.10), reckoning in the unit (Section 13.7.3), and the state-contingency of differential value (Section 13.3.11). Ratings that go up are self-congratulation; ratings that identify the unmeasured bridges carry information.

The rest of this Part defends links 7 through 9, where the monetary argument actually lives. Parts III and IV defend links 2 through 6 by building and measuring the stack that supplies them. Parts V and VI make every link falsifiable.

8.2 Topological Scarcity and the Pressure–Capacity Corridor

Topological Scarcity Lemma

A nominal stock of resources, capital, or capacity does not secure a monetary service. The relevant quantity is the stress-deliverable flow through the narrowest non-substitutable edge connecting input to holder utility.

Aggregate abundance and local scarcity can therefore coexist without contradiction. Crude is not diesel when refining is the minimum cut; savings are not duration absorption when no balance sheet can hold DV01; installed accelerators are not verified service when power, cooling, firmware, connectivity, verification, or settlement fails. The topology—the transformation path, location, actor, and horizon—determines the service.

The physical-conversion example is adapted selectively from Doomberg's *Endangered Specious* [Doomberg 2026b]; the duration-absorption example from Michael Green's *Anchors Aweigh, My Boys* [Green 2026b]. Both are analyst commentary used for mechanism rather than quantities. The lemma and its application to verified service are this thesis's own.

Link	Strength	Why
Utility demand	Strong	Structural drivers are already observable; the delegation mechanism (Chapter 7) is named, its instrument unbuilt
Standardized work	**Medium**	PoUW constructions are a research direction, not a shipping monetary base; delegation is a standardization pressure, not a construction
Receipts	Strong	Structurally well specified; VERIFYPRICE makes it measurable
Stress-deliverable service	**Medium**	Scenario max-flow, binding constraints, and common-cause data are specified but not yet demonstrated at scale
Markets	**Medium**	Market microstructure for proof and compute capacity is the least developed part of the design
Fee flows to asset	Strong	Formalized as a lemma with public boards—but establishes *cash-flow accrual*, not monetary premium (Section 13.3.9)
Non-bypassability	**Medium**	Written as a binary; the economics requires a magnitude (Δ) that is nowhere estimated (Section 13.3.4)
Native holder anchor	**Medium**	Transactional, collateral, and wrapper demand do not establish countercyclical loss-bearing capacity; nor does a cohort that cannot be classified by principal
SoV premium	**Medium (conditional)**	Nine requirements, eighteen red lines; premium derives from link 9 through three mechanisms — pledgeability, denomination, holder-side flow — not from fee accrual (Section 13.3.10, Section 13.7.3, Section 13.3.11)

Table 8.1: Link-by-link strength of the conditional chain. Six links remain Medium or Medium (conditional): standardized work and market formation are engineering-maturity risks; stress-deliverable service is a physical-topology and common-cause risk; non-bypassability is an unmeasured magnitude; the native holder anchor is an unproven risk warehouse; and monetary premium remains conditional on the entire chain.

Let R denote regime pressure, $\Delta(R)$ the holder's differential value over the best institutional substitute, $A_{\text{full-stack}}(R)$ the availability of the complete service path under that pressure, and $H(R)$ the holder's practical ability to self-custody, transact, prove, and exit. The monetary service scales with the product of the three, on the product-form premise

$$\Pi_{\text{monetary}}(R) \;\propto\; \Delta(R) \times A_{\text{full-stack}}(R) \times H(R).$$

This is the **Pressure–Capacity Corridor**. Its likely shape is an inverted U, not a crisis escalator. At low pressure, courts, custodians, and indemnitors work, so Δ is thin. At moderate pressure, substitutes weaken while the open stack and holder remain operational, so differential value can rise. At extreme physical or political pressure, power, networks, hardware, liquidity, custody, or exit may fail and the service falls with them. More crisis does not imply more premium.

$H(R)$ has an agent form. A software agent's practical ability to self-custody, transact, prove, and exit is bounded by its operator's: a fleet instance inherits one jurisdiction's pressure all at once, while independently operated agents inherit it operator by operator. That is the coercibility split of Section 7.4 in one symbol, and it is why a holder base made of agents must be read by operator class before $H(R)$ is credited to it.

8.3 The Triad Coherence Test

Privacy settlement, portable proof, and verified compute are different goods with different users, bottlenecks, collateral needs, political risks, and duration warehouses. A common asset is justified only if shared security, settlement, liquidity, and transaction-cost benefits exceed cross-subsidy, governance, and wrong-way-risk costs.

For each architecture ask: Do the services share a security budget and users? Does demand for one improve the economics of the others? Are bottlenecks complementary or correlated? Does common settlement reduce transaction costs more than it creates cross-subsidy? Can one constitution credibly govern all three? Does one service dominate issuance, fees, or political risk? Architecture A is a candidate, not the thesis.

8.4 Monetary Objects and Value Capture

The thesis claims that triad capacity can earn a "durable store-of-value premium." For that claim to be testable, four questions must be answered without poetry:

1. **What exactly is the asset?**
2. **What do users pay in?**
3. **How does holding capture value?**

Architecture	Strength	Burden of proof
A: One native monetary asset	Strongest direct value capture and shared liquidity	Greatest reflexivity, governance scope, and cross-domain collateral risk
B: Neutral reserve plus service-specific credits	Clean separation between reserve money and typed service claims	Weaker new-token capture; requires interoperable settlement and independent credit markets
C: Shared settlement plus modular domain collateral	Common payment rail with service-specific risk containment	More operational complexity; collateral fragmentation and bridge governance must remain legible

Table 8.2: Triad architectures under test; none is pre-selected.

4. **Why does value not leak entirely to operators?**

Question 1: What is the asset? The triad stack can issue several kinds of holdable instruments:

Instrument	What It Represents	Scarce?	Transfer?	SoV?
Base Asset	Conditional bearer monetary candidate; fee and settlement unit	Yes	Yes	Candidate only
Work Credits	Typed capacity or service claims, bounded by deliverable output	Yes	Yes	Not presumed
FCR/FER/PIDL artifacts	Evidence of physical conditions or completed work	No	Yes	No
Project Notes	Explicit duration-bearing infrastructure credit	Yes	Yes	No
LP/Staking	Rights to fee share from corridors, pools, validators	Yes	Sometimes	Derivative
Capacity Vouchers	Prepaid access at fixed rates	No	Yes	No

Table 8.3: Instrument hierarchy. Utility, evidence, credit, derivatives, and the conditional monetary candidate remain distinct.

The **base asset** is only a conditional monetary candidate: the unit in which fees may be paid, burns may occur, and collateral may be posted. Work Credits can be valuable, transferable, workload-specific, location-specific, SLA-specific, and temporally perishable without becoming savings instruments. FCRs, FERs, and PIDL receipts are evidence; project notes

warehouse duration and bear default risk; LP and staking positions are derivatives or operating claims.

A terminology contract. This thesis uses two distinct words that are easily confused, and the distinction carries real analytical weight:

- **Native instruments** are protocol-internal claims: the base asset, typed Work Credits, staking and LP positions, corridor claims, capacity vouchers, and evidence artifacts. They touch the protocol, but only the base asset is evaluated as a monetary candidate.
- **External financial wrappers** are conventional market products written *on* the asset: spot ETFs, exchange-traded products, treasury companies, custodial balances, futures, options, swaps, and leveraged or inverse ETPs. They do not touch the protocol. Holding them exercises nothing.

Both are sometimes called "wrappers" in casual usage. From here forward, "wrapper" without qualification means an *external financial wrapper*. The full object hierarchy runs from physical capacity out to recursive financial claims, and monetary treatment differs sharply along it.

Object layer	Example	Touches protocol?	Path dependence	Monetary treatment
Service capacity	Privacy, proof, compute throughput	Yes	No	Non-monetary input measured by DVC
Typed service claim	Work Credit / capacity voucher	Yes	Medium	Not presumed money or SoV
Native monetary object	Base asset	Yes	Low	Conditional SoV candidate
Native derivative	Staking or LP shares	Yes	Medium	Secondary claim
Custodial claim	Exchange balance, spot ETF share	Indirectly	Low–medium	Price exposure, not full monetary function
Synthetic wrapper	Futures, swaps, options	No	Medium–high	Financial exposure only
Daily-reset leveraged wrapper	$2\times$ / $3\times$ ETP	No	Extreme	Trading product; explicitly not SoV
Recursive wrapper	Leveraged product on another wrapper	No	Extreme, nonlinear	Systemic-risk instrument

Table 8.4: Object hierarchy from base capacity to recursive financial claims. Monetary properties do not survive the descent: a daily-reset leveraged wrapper cannot inherit the underlying asset's store-of-value status, because its return depends on the path, not merely the endpoint (Appendix H.9).

How to read Questions 2–4. The answers below are stated in the reference-design register: they specify *routing rules* — where fees go, what must be locked, what may not be inflated away. They are design commitments, not findings, and two things they deliberately do not do. They do not establish that the routed fee is *large*: that is the level question, taken up in Section 13.3, where the sustainable fee is bounded by differential value Δ rather than by anything in this section. And they do not establish that any of this is monetary premium rather than a cash-flow claim: Section 13.3.9 shows why that reading fails, and Section 13.3.11 relocates the monetary claim to a holder-side service flow. The percentages below are illustrative of a design space, not measurements of a market.

Question 2: What do users pay in? In the reference design, all triad services are *quoted and settled* in the **base asset**:

- **Proofs:** paid to provers in the base asset; a published share retired.
- **Privacy settlement:** corridor fees paid in the base asset; a published share retired.
- **Verified compute:** inference and MatMul paid in the base asset; a published share retired.
- **Staking/collateral:** provers, routers, and LPs lock the base asset to participate.

This is a *routing rule*, not a demand theorem. It creates transactional demand only to the extent that users cannot route around it (Value-Capture Condition 1; Red Line 6), and even where they cannot, the demand it creates is for a *medium of settlement* held for seconds, not for a store of value held through a cycle. Section 13.3 bounds what that toll can be, and Section 13.5.1 names the incumbent that already routes around it. Nothing in this list is evidence of moneyness.

Question 3: How does holding capture value? Fee value reaches holders through four routing channels. Each is a cash-flow or lockup mechanism; none is monetary premium, and the thesis says so at Section 13.3.9:

1. **Fee retirement (burn)** (30–50% of fees): a pro-rata buyback. It reduces *net* issuance only when retired fees exceed scheduled issuance, which at early-phase fee coverage they will not; it is published as its own series and read as such.
2. **Staker yield** (20–40% of fees): income for holders who stake — a dividend on a productive claim.
3. **Operator share** (20–40% of fees): incentive for provers, routers, and LPs.
4. **Collateral requirements** (10–20% of capacity value): immobilizes float while operators remain solvent and choose to keep operating; Section 13.3.12 shows why it is a slope effect on price impact rather than a valuation floor or proportional accrual, and bounds the fixed dollar quantity it can lock.

The first three channels are a **partition of one fee pool**, not three independent receipts: burn share, staker yield, and operator share must sum to 100% of every fee collected. The

bands above are alternative interior choices within that partition, not independent ranges, and any point allocation — such as Design A's 40/30/30 split (Section 8.5) — must close the identity. Reading the bands as additive would sum the maxima to 130% of the pool, which is not an allocation the design can make. The **reference design** this document actually works from is fixed once, in Section 25.6: 70% to capacity providers and the stakers who delegate to them, 20% retired, and 10% to an assurance budget that funds Layer 0 sampling, audits, and security review. Its 20% burn sits below the band above deliberately, as the harder case for the value-capture claim (Chapter 22), and every worked example in Parts III–V uses that point; Design A's split is an alternative inside the bands, not a second reference.

Question 4: Why does value not leak entirely to operators? The "utility token trap" occurs when operators capture all economic value while token holders merely provide exit liquidity. The routing rules are designed so that a positive fee, *if one exists*, does not leak entirely to operators:

- **Fee retirement:** 30–50% of fees are retired pro rata to all holders — a buyback, with the accrual (not moneyness) reading Section 13.3.9 gives it.
- **Staking yields:** Passive holders can stake to earn fee share without operating infrastructure.
- **Collateral requirements:** Operators must hold significant base-asset collateral, aligning their interests with holders.
- **Governance rights:** Holders vote on fee splits, issuance changes, and protocol upgrades, within the constitutional constraints of Chapter 25.
- **Issuance constraints:** Base-asset issuance follows the constitutional schedule and is independent of Work Credit dynamics; it is *Work Credits* that are sold only against committed, reserved capacity, with any sale beyond deliverable capacity rejected and visible in dashboards.

Net effect, routing only: these rules route a fee of size $f \cdot Q$ to retirement, staking, and operator share. Whether f can be positive at all is a function of differential value Δ, not of routing; whether the routed revenue constitutes monetary premium is a separate question this Part takes up in Section 13.3 and answers in the negative at Section 13.3.9.

8.5 Three Reference Designs

To make the thesis testable, it commits to three concrete designs. Implementations may vary, but at least one must be viable for the SoV claim to hold.

Design A: Base Token as SoV (Primary)

- **Issuance:** Fixed schedule with halvings (like BTC), set ex ante.
- **Fee medium:** All triad services priced in base token.

- **Burns:** 40% of fees permanently destroyed (a point choice inside the 30–50% reference band).
- **Staking:** 30% of fees to stakers; 10–15% collateral requirement.
- **Operators:** 30% of fees to provers/routers/LPs.
- **Governance:** Token-weighted voting on parameters.

SoV properties: Credible scarcity (capped + burns), native demand (fees), duration-neutral (no coupons), cheap verification (VerifyPrice dashboards).
Risk: If demand stalls, burns decline and scarcity weakens.

Design B: Work Credits as Capacity Vouchers (Hedge Instrument)

- **Issuance:** Sold against committed capacity (FERs + FCRs), reserved atomically; outstanding claims bounded by stress-adjusted DVC; no fixed cap.
- **Redemption:** Extinguished by verified delivery at the contracted tier and rate; a priority tier is a term of the contract, paid for when the credit is bought.
- **Expiry:** Each credit carries its delivery window; unexercised credits expire or refund on stated terms, which prevents hoarding and bank-run dynamics.
- **Transferable:** Yes, but primarily used for cost hedging, not long-term savings.

Use case: Enterprises hedge against compute cost spikes; treasuries lock in future settlement capacity.
Not a SoV: Supply expands with capacity; expiry prevents indefinite accumulation. This is infrastructure hedging, not a store of value.

Design C: Triad Index Token (SoV Candidate)

- **Backing:** Diversified revenue streams across privacy corridors, proof pools, and compute markets.
- **Value capture:** Protocol revenue used for periodic buyback-and-burn or dividend distribution.
- **Issuance:** Fixed supply; no new minting after genesis.
- **Governance:** Index holders vote on revenue allocation and rebalancing.

SoV properties: Diversified exposure to triad demand; fixed supply; yield via buybacks or dividends.
Risk: Concentration in specific corridors or proof markets; governance capture.

Which design does the thesis endorse? None by default. These instrument designs sit inside the broader architecture comparison of Section 8.3. A fixed-supply base token may be tested as the monetary candidate; capacity vouchers remain service hedges; an index token is a revenue-bearing security whose buybacks or dividends do not establish moneyness. The text uses

"base asset" for the monetary candidate and "Work Credit" only for a typed capacity or service claim.

What actually differs, and who loses. The three designs are not three shades of the same bet; each one moves the risk onto a different balance sheet. Design A concentrates it in the base asset: if demand stalls, burns decline, and the scarcity story thins with it — the holder discovers that a burn schedule is a claim on volume, not a property of the token. Design B pushes the risk onto the service buyer: capacity vouchers hedge compute cost beautifully until capacity is scarce in exactly the state the buyer hedged for, at which point delivery terms rather than price are the binding constraint, and the hedge is only as good as the stressed path behind it. Design C hands the risk to governance and to the revenue mix: an index token is a claim on whatever the protocol happens to earn, so concentration in particular corridors or proof markets is not a footnote to the thesis but the thesis, and buybacks do not convert a revenue stream into money (Section 13.3.9). Watch the same three numbers in all three cases — fee share actually paid in the native asset, net supply after burns, and the share of exposure sitting in wrappers (Section 13.6) — because Section 13.3 takes up the question these boxes deliberately do not answer: not where the fee is routed, but whether there is a fee to route.

Chapter 9

Privacy as Private Money

Cash used to be **default private money**: anonymous, bearer, final on receipt. In a world of KYC'd banks, programmable payments, and networked surveillance, that role has decayed. At the same time, the repression playbook (negative real rates + capital controls) makes **bearer-like savings** economically necessary, not ideologically optional.

"Privacy" in this thesis is not romantic opacity; it is **the ability to hold and move value without chokepoints, plus the option to disclose on your own terms.**

Interiority has a history. The capacity this chapter argues for has a five-hundred-year arc behind it, and the arc is worth naming because it reframes what is being defended. The private interior self is not a natural constant that technology is eroding; it is, on the account of the orality-literacy literature [Ong 1982], substantially a *product* of a technology: sustained silent reading, a practice individuals had to acquire, that only became mass practice with print. Interiority was an externality of a medium, and it has since been mistaken for the default state of persons. On that reading, surveillance does not violate a natural privacy so much as withdraw a subsidy that a particular medium had been paying, quietly, for centuries. The implication for this chapter is not that the case for private settlement strengthens — the repression economics stand on their own — but that the loss is larger than it looks: what is being unbuilt is not an amenity but a specific, technology-dependent form of the person that took centuries to build and that no one planned. Defending it is therefore also not an amenity.

Concretely:

Bearer-like holding: Keys, not accounts, define control. Custodians may exist, but custody is an *optional service*, not a mandatory chokepoint.

Non-custodial settlement: Cross-asset flows (BTC↔XMR today; BTC↔shielded-ZEC when a corridor exists, Section 23.3) execute as atomic swaps or corridor protocols; neither side needs to trust a centralized intermediary.

Auditable by consent: Viewing keys and structured receipts allow specific flows to be disclosed to auditors without exposing the entire graph.

When these properties hold, **private settlement capacity itself** starts to behave like a monetary asset:

- A treasurer facing capital controls is not just asking "what's the yield?" but rather, "*can I*

still get value to my people next year if on-/off-ramps are throttled?"

- A dissident journalist or NGO cares less about upside and more about *"will this still be here and spendable if my local banks freeze?"*

> **Blockchain may be waiting for its SSL moment.** The analogy is suggestive: in 1994, putting credit card information on the Web seemed reckless until SSL made encrypted commerce viable. E-commerce grew into a multi-trillion-dollar industry once confidentiality became default. Today, most public blockchains are "public by default" the way HTTP was—every transaction visible, every address linkable. Privacy is plausibly the bottleneck for mass institutional adoption.

The institutional version is straightforward: no enterprise wants payroll, vendor rates, or treasury operations visible by default. The unlock is one-click, non-custodial BTC↔XMR (and BTC↔shielded-ZEC once that corridor exists) with refund-safe UX and clear settlement analytics—privacy as a product, not a promise.

In that environment:

- The **rails** (privacy corridors, shielded pools) earn fees for providing unseizable, auditable settlement capacity.
- The **claims** on those rails (e.g., corridor LP positions and Work Credits) are service, operating, or derivative claims. They may be held, but are not presumed stores of value.

You can think of **Private Money** as:

> *"Rights to future, censorship-resistant, auditable settlement capacity."*

The SoV properties from Chapter 4 map naturally:

- **Credible scarcity.** Capacity is constrained by bandwidth, cryptographic verification costs, and capital at risk.
- **Cheap public verification.** Anyone can verify that a transaction was correctly formed, swapped, or refunded.
- **Censorship-resistance & portability.** Flows ride over non-custodial corridors; exit options span multiple assets and jurisdictions.
- **Neutrality & permissionlessness.** Adversarially diverse relays, routers, and LPs; public metrics on concentration.
- **Native demand.** Demand is created by repression itself.
- **Lawful privacy.** Viewing keys + PIDL receipts mean regulated entities can prove compliance without deanonymizing entire networks.
- **Duration-neutrality.** Claims participate in fee flows and scarcity premiums, not fixed coupons.
- **Asset-level value capture.** Fees, collateral, and settlement are denominated in the base asset, so consuming the service touches the asset rather than bypassing it (Section 13.5).

- **Agency preservation.** Selective disclosure and non-custodial settlement equip the holder rather than managing them (Item 9).

The ninth requirement carries a capability reading the eighth does not, and the distinction matters because a ledger can satisfy the formal properties while failing the practical one. Formal access — a valid balance, a working key, a legal right to transact — is not the same as capability: command over the necessities of an independent life. A perfectly accurate ledger can faithfully preserve an unequal starting state; making claims harder to alter does not broaden who can acquire them (Section 34.3.6). The requirement as enforced (Section 26.9) therefore reads on both layers: the Agency Preservation Board tracks non-custodial usage share and disclosure scope — the formal layer — and the participation telemetry of Section 3.6 tracks whether users can in practice acquire and hold meaningful claims through wages rather than inheritance or incumbency — the capability layer. A stack that passes the first and fails the second has hardened the incumbent distribution more efficiently than fiat did.

On a balance sheet, Private Money can coexist with BTC and fiat, but it plays a different role:

- **BTC**: thermodynamic base, global risk asset, macro hedge.
- **Fiat**: near-term unit of account, legal tender, credit medium.
- **Private Money**: repression-hedged rail, a way to ensure you can still pay and be paid without being fully seen.

Chapter 10

Proofs as Attestation Money

If Privacy is the hull, **Proofs** are the **receipts**.

Proofs Are Not Truth

Proofs are bounded attestations, not truth. A proof can show that a computation followed a circuit, that a file descends from a signed provenance chain, or that a settlement rule executed correctly. It cannot prove that a camera was pointed at what matters, that an event was representative, that a witness is honest, or that an AI output is substantively correct. Throughout this chapter, "proof" means "cryptographic attestation under stated assumptions," not semantic or social truth.

In a world of synthetic media, platform moderation, and "liar's dividend," the scarce thing is no longer *content* but **trustworthy provenance and computation history**. Proofs are the mechanism layer that answers, cheaply and publicly:

> *"Did this actually come from where it claims, and did the computation actually run as stated?"*

Clarification: Proofs vs. Proof-Backed Instruments Proofs themselves are *infinitely replicable outputs* once generated. A ZK proof can be copied and verified anywhere. On its own, a proof is not a scarce bearer asset—it is an attestation.

What *is* scarce is:

- **The ability to produce valid proofs at scale** (requires compute, hardware, energy).
- **Rights embedded in instruments** that reference proofs (Work Credits, staking positions, capacity claims).

Proofs therefore play three distinct roles in the monetary stack:

1. **Proofs as a commodity market:** Standardized attestations priced by VerifyPrice. You buy proofs the way you buy bandwidth—as a priced input to operations.
2. **Proofs as collateral enablers:** Proofs make *other* contracts collateralizable. An inference SLA backed by proofs of correct execution can be used as collateral because the proof makes default detectable.
3. **Proof-backed service instruments:** Work Credits tied to verified proof workloads are typed service or capacity claims whose integrity is enforced by proofs. They are not presumed

money.

"Attestation Money" survives as an analytical lens or legacy label for proof-demand economics. The instruments beneath that lens are proof-backed service claims; any monetary candidacy belongs to the separate base asset.

From a monetary perspective, two features matter:

1. **Proofs travel:** A PIDL receipt can be moved across systems, archived, or collateralized. It outlives any single platform's UX, TOS, or reputation.
2. **Proofs can be priced and standardized:** Once you know the VerifyPrice of a workload, you can treat "valid proof of workload W" as a commodity unit.

This gives rise to an **Attestation Money analytical lens**:

- Media platforms, insurers, and courts demand proofs of origin, editing history, and custody.
- AI services demand proofs that a model of a given hash ran on given inputs with given bounds.
- Enterprises demand proofs that compliance computations actually ran.

Initially, proofs are purchased as **opex** ("we pay per proof"). Over time, markets will create **claims on future proof capacity**:

- Reservations or futures on proof-of-provenance capacity for a media network.
- Proof pool shares that entitle holders to a portion of fees from high-value workloads.
- Work Credits sold against committed proof capacity that meets certain VerifyPrice and SLO thresholds.

Under the legacy **Proof/Attestation Money** analytical lens, these instruments are typed claims on proof capacity, not monetary instruments. Work Credits remain service claims; only the separate base asset may become a monetary candidate if the full conditional chain passes:

> *"Rights to future, standardized, verifiable attestations about data and computation."*

From the triad perspective:

- Privacy protects *who* is involved.
- Proofs protect *what* actually happened.
- Compute powers *how* we arrive at outputs.

The "Attestation Money" lens describes demand arising where **verification, not vibe, mediates trust**. Claims on proof capacity remain service instruments; they are not what this thesis asks holders to treat as money.

Chapter 11

Compute Through the "AI Money" Lens

Compute has always been an economic input: first as muscle, then as steam and electricity, now as FLOPs. The AI boom made this explicit: GPUs, TPUs, and datacenters are priced like oil fields.

Raw compute is not even a standardized service claim:

- Most compute markets are opaque capacity rentals (cloud contracts, colo deals).
- There is no standardized unit of **verified work**; only hours and instance types.
- There is no cheap, public way to verify that a claimed computation actually ran correctly.

Research Maturity Caveat

Verified inference remains technically immature. Only Tier A cryptographic verification (full ZKML) should support high-assurance capacity issuance without an additional trust haircut. Probabilistic or TEE-backed inference verification (Tiers B and C) receives haircuts or lower-assurance service status. The service-value claim weakens if verified compute trades at no durable premium to ordinary cloud compute; that result would say nothing by itself about monetary premium. See Chapter 24 for the tier taxonomy.

The "AI Money" analytical lens reframes verified-compute demand by insisting on:

- **Canonical workloads:** Matrix multiplications, FFTs, and core model primitives.
- **Proofs of useful work (PoUW):** Miners/provers earn rewards for producing proofs that these workloads ran correctly, with verification asymmetry.
- **Verification economics:** VerifyPrice turns "one unit of verified MatMul at dimension n" into something that can be priced and traded.

In that context, **verified compute capacity** becomes a standardized service commodity:

> *"Rights to future, standardized, cheaply verifiable units of useful compute (verified FLOPs)."*

11.1 Compute Deflation and What Remains Scarce

A skeptical reader will object: *"Compute gets cheaper every year. Moore's Law drives raw FLOP cost down. How can rights to compute retain service value when the underlying commodity deflates?"*

This is a serious objection. Here is the response:

What deflates:

- Raw FLOPs per dollar (secular decline).
- Cost of unverified, permissioned compute from hyperscalers.
- Commodity inference for non-sensitive workloads.
- **The price of installed capacity after a clearing event.** Compute built on financing written against utilization assumptions that did not arrive can change hands at a fraction of its capital cost with the silicon unchanged — financial deflation [Perez 2002] on top of technological deflation [Booth 2020]. See Section 13.3.2 for the effect on the fee base.

What remains scarce:

- **Verified FLOPs with receipts:** Compute where correctness is cryptographically proven, not vendor-asserted.
- **Policy-constrained capacity:** Compute that runs under specific jurisdictional, privacy, or compliance constraints.
- **Censorship-resistant access:** Compute that cannot be denied by TOS changes, sanctions, or platform bans.
- **Priority access under congestion:** During demand spikes, priority slots are scarce even if raw capacity is abundant.
- **Specific hardware profiles:** Compute on L0-C or L0-D grade hardware may remain scarce even as closed alternatives proliferate.
- **Receipts for delegation across trust perimeters:** FLOPs deflate; the boundaries between operators that a decomposed task must cross do not, and cheaper compute makes finer decomposition cheaper (Section 7.2). Whether those hops cross perimeters at all is an open measurement, not a forecast.

The instrument should reference capacity share, not "one timeless FLOP": Under the "AI Money" lens, the relevant service claim is not "rights to 1 FLOP forever." It is:

> *"Rights to X% of verified compute throughput under SLA Y on hardware profile Z."*

This is analogous to how oil futures reference "barrels of WTI crude delivered at Cushing"—not "energy" in the abstract. The specificity (verification, SLA, profile) is what creates persistent scarcity.

Telemetry that detects deflation risk: If raw compute deflation outpaces the scarcity premium of verification/censorship-resistance, the following metrics will signal it:

- VerifyPrice for verified FLOPs converges toward unverified market rates.

- Capacity utilization on verified networks falls below thresholds.
- Fee revenue from compute workloads declines as a share of total fees.

These are observable. The thesis predicts that a verification service premium may persist because demand for trustworthy AI is structural and the compliance/sovereignty premium may grow as AI becomes more consequential. If telemetry shows otherwise, the compute service case weakens. Neither outcome determines whether the separate base asset earns monetary premium.

Verified-compute service instruments can include:

- **Work Credits**: typed claims on reserved verified-compute capacity, owed by a named provider and redeemable or transferable for specified future workloads.
- **Compute futures**: for specific workloads (e.g., "X verified inferences at model M with SLO Y").
- **Stake in proof factories**: whose entire business is converting energy + silicon into verified FLOPs with transparent VerifyPrice.

Why can these be valuable service claims rather than arbitrary utility tokens? Because:

- **Demand is global and secular:** As long as AI is an input to value creation, there is demand for FLOPs.
- **Verification is public:** Anyone can check that a Work Credit corresponds to a valid proof and a stated DVC envelope.
- **Terms are explicit:** Workload, tier, delivery window, redemption, and expiry determine value rather than an implied monetary promise.

Chapter 12

Work Credits: Reserved Service Claims

Every era that mechanizes work eventually writes claims on it. Grain receipts became banking; warehouse receipts became money-market funds; and in each case the instrument was born useful and lived or died on one question — did the claim stay tethered to the thing in the warehouse? The tether is everything. A receipt that outruns its warehouse is a promissory note wearing a receipt's reputation, and the bill arrives the day everyone redeems at once. This chapter builds the tether before it builds the claim.

The triad supplies three capacities. Making them tradable without confusing service with money requires a **unit of account for work**: something that binds energy, hardware, delivery terms, and verification into a transferable typed claim.

Call these **Work Credits (WC)**.

Informally:

> *A Work Credit is a typed, prepaid claim on reserved future service, owed by a specified provider or provider pool. Its issuance requires an admissible capacity commitment and an atomic reservation against outstanding obligations. Verified fulfillment extinguishes the claim and produces a receipt, which is evidence and creates no further obligation.*

Key distinction: Work Credits are *energy-priced, not energy-pegged*. They are denominated in work, not in kWh. Energy enters through Facility Energy Receipts (FERs) and VerifyPrice, not through a one-dimensional peg. See Section 12.7 for the full treatment.

Clarification. Throughout this section, "Work Credits" are generic, notional service claims, not product branding and not the base asset. "AI Money" and "proof money" survive only as historical shorthand for service-instrument design spaces; the analytical objects are typed compute, proof, and settlement claims. That separation — claim versus money — is stated once here and defended twice, at Section 12.4 and Section 12.8; it is not repeated as a caveat at every mention.

Work Credits can be described along a few axes:

12.1 Definition: Work Receipts vs. Work Credits

A common source of confusion is conflating "proof that work was done" with "transferable claim on future capacity." These are different financial objects. They split cleanly:

Work Receipt (WR)

A **Work Receipt** is a PIDL artifact proving that a specific unit of work was completed under attested conditions.

- **Content:** Claim hash, proof hash, workload ID, SLA tier, timestamps, hardware profile, prover signature.
- **Properties:** Copyable, verifiable by anyone, *not scarce*.
- **Analogy:** A receipt from a completed transaction. It proves the past but confers no future rights.

Work Receipts are *not* money. They are evidence.

Work Credit (WC)

A **Work Credit** is a transferable, typed claim on reserved future service, issued against an admissible capacity commitment and extinguished by the Work Receipt that verified delivery produces.

- **Issuance:** Sold when (a) a named provider or constituted pool has committed admissible capacity for the workload, tier, and window, and (b) that capacity is reserved against every outstanding claim on it in the same atomic step.
- **Properties:** Scarce (outstanding claims bounded by reserved capacity), transferable, fungible within workload, tier, and window.
- **Rights:** The service terms written in the claim: workload, quantity, tier, delivery window, substitution terms, and remedies.

Instrument	Rights Conveyed	Scarce?	SoV?
Base asset	Bearer fee, settlement, and collateral unit under a constitutional issuance rule. No claim on historical or future work; holding it confers no priority, capacity, price, or queue right.	Yes	Conditional candidate only
Work Credit / WC-Voucher	Typed prepaid claim on reserved proofs, compute, or settlement capacity, owed by a named provider or pool at stated workload, tier, place, and delivery terms.	Bounded by reserved capacity (DVC)	No (service/hedge instrument)

Base asset and Work Credit separation: **Hierarchy rule:** The base asset is the only object evaluated for monetary premium, and even then only conditionally. Work Credits are useful for service procurement and cost hedging and are explicitly not pitched as stores of value.

12.1.1 The Work Credit Lifecycle

A Work Credit lives through four steps, and one of them is indivisible:

$$\text{commit capacity} \;\rightarrow\; \boxed{\text{sell} + \text{reserve}} \;\rightarrow\; \text{deliver} + \text{verify} \;\rightarrow\; \text{extinguish} + \text{receipt}.$$

There is no interval in which a valid sold claim exists while its capacity is unreserved. Four commitments make the lifecycle a contract rather than a slogan.

- **The obligor is named.** Each credit binds a provider or an expressly constituted provider pool. "The network" is not an answer unless the network's collective obligation, funding, and loss allocation are themselves defined.
- **The service is named.** Workload, quantity, assurance tier, delivery window, location or substitution terms, and remedies are written in the credit. A substitution right states what may replace what.
- **Reservation is against everything outstanding.** The constraint is not how many credits were sold this epoch but whether every outstanding claim competing for the same future capacity remains feasible. Shared power, hardware, and delivery windows cannot be reserved independently for incompatible workloads, which is why the reservation limit is computed jointly (Section 17.9).
- **Failure resolves by terms.** Verified delivery retires the claim. Cancellation, expiry, default, and refund follow stated terms, and a failed delivery becomes neither an open-ended promise nor an automatic claim on the base asset.

Delivered Verified Capacity supports the reservation limit; it does not guarantee performance. In the reference design, payment sits in the fulfilment and escrow arrangement of the claim until delivery rather than quietly financing construction. Where advance proceeds do finance capacity not yet built, that exposure is project finance, named and carried separately (Section 3.3), and not concealed inside the ordinary Work Credit. The receipt that delivery produces is evidence of fulfilment and creates no further service obligation, which is the line between this section's two objects.

Issuance mechanics for Work Credits: For a canonical workload W (e.g., "MatMul of size n with error bound ε," "provenance proof for content type C," "corridor settlement of size S with anonymity set $\geq A$"), define:

- $p(W)$: production cost (energy + hardware amortization + opex) to generate one unit and produce a valid proof.

- $v(W)$: verification cost.
- $r(W) = v(W)/p(W)$: verification asymmetry.

A **Work Credit of type W, tier T** is sold only when:

- A named provider or pool commits admissible capacity for W at T over a stated window, with FER and FCR evidence behind it.
- **Telemetry confirms** that VerifyPrice(W,T) and other SLOs (latency, failure rate, decentralization) are within bounds for that provider.
- The sale and the reservation happen together, and outstanding claims on the same capacity stay within the **issuance envelope** (see Section 12.6).

The credit represents only the service rights stated in its terms. A receipt attests historical work and a credit promises reserved future service; neither becomes the fee/staking medium or inherits the base asset's conditional monetary candidacy.

12.2 Energy Anchoring

Like PoW, Work Credits are ultimately **energy-anchored**:

- The cost of honoring one more credit is bounded below by the energy and hardware required to produce work that passes the verification threshold.
- Where SHA-256 PoW has no buyer for its work, this work has one: it powers privacy settlements, proofs of provenance, and AI computation.

This anchoring gives Work Credits:

- **Credible scarcity:** A credit cannot be sold without an admissible capacity commitment, and it cannot be extinguished without real resources spent producing proofs and settling flows.
- **Economic meaning:** One credit corresponds to a service somebody was willing to pay for (anonymized payrolls, authentic media, verified inference).

And the cost of the same property. The buyer that gives a Work Credit its economic meaning is also an attack surface the hash puzzle does not have. Nobody buys a SHA-256 output, so there is no buyer of the work to subsidize, mandate, or sanction — though miners themselves can be; demand for verified inference can be all three. The energy anchor is therefore a service-cost and provenance anchor, not a monetary one. Bitcoin's monetary objectivity and a Work Credit's link to real activity are different properties (Section 34.1.1).

12.2.1 Layer 0 Maturity and Economic Consequences

Work Credits are only as trustworthy as the hardware that produces them. A world where all proving runs on opaque, vendor-controlled hardware is different from one with diversified, partially open designs.

Chapter 17 (Part III) defines a **Layer 0 maturity ladder** with grades L0-A through L0-D. This section previews how those grades affect Work Credit economics:

L0 Grade	Hardware Trust Level	Economic Treatment
L0-A	Closed hardware + attestation + multi-party audits + diversity	WC accepted at full value if diversity thresholds met; risk premium priced into fee schedules
L0-B	Lot sampling + imaging + bounded side-channel budgets	WC accepted at full value; lower risk premium
L0-C	Partial open (open RTL for critical components)	WC may qualify for "open-profile" premium or priority tiers
L0-D	Fully open designs	Highest trust tier; may command premium in markets that value sovereignty

How this affects issuance and pricing:

1. **Tiered issuance caps:** WC reserved on L0-A profiles may face stricter issuance limits than L0-C/D profiles. This prevents the system from becoming dependent on opaque hardware.
2. **Market pricing:** WC from different L0 grades can trade at different prices if markets distinguish them. Treasuries and institutions may pay premiums for L0-C/D-backed capacity.
3. **Risk disclosure:** Every Work Receipt includes the hardware profile (HID) used. Aggregated telemetry shows what fraction of total WC is backed by each grade.
4. **Deprecation impact:** If an L0-A profile is compromised (TEE backdoor discovered), WC reserved on that profile can be quarantined, discounted, moved to substitute capacity under their terms, refunded, or excluded from certain uses. This is the "bond downgrade" analog.

Why this matters for service-claim quality: If 90% of Work Credits are reserved on L0-A hardware and a major TEE is compromised, claim quality and DVC take a hit—not because the cryptography failed, but because the service path depended on a common hardware assumption.

By making L0 grade explicit and tying it to economic consequences, the system:

- Creates incentives to invest in open hardware.
- Bounds the impact of hardware compromises.
- Gives users and allocators the information to price risk appropriately.

This is covered in detail in Section 17.4 (Part III). The key point for Part II: **hardware trust is not binary; it is graded, measured, and priced** — and because it is priced, the service claim's quality is legible without anyone having to take the word "typed" on faith.

12.3 Robots, AI, and the Demand for Work Money

In a robot- and AI-heavy economy the demand story for Work Credits becomes almost embarrassingly straightforward.

Consider a warehouse or factory where most of the physical activity is carried out by robots, and most of the planning and oversight is handled by models. The day-to-day budget splits into:

- Energy to power robots and data centers.
- Compute to run models and generate proofs.
- Privacy and settlement to pay workers, suppliers, tax authorities, and investors without leaking trade secrets or exposing everyone's graph.

Each of those budgets expresses itself as a recurring need for receipts:

- Proof of correct inference for high-stakes decisions.
- Proof of compliance and risk calculations for regulators and insurers.
- Privately settled wages and vendor payments with lawful audit trails.

Today that flow is mediated through cloud bills, payroll files, bank wires, and audit PDFs. In the stack described here, it can be mediated through Work Credits and receipts. The robots and services are paid in typed claims; in exchange they produce receipts that can be verified cheaply and settled over privacy rails. High turnover can make the claims useful working balances. It does not establish reserve demand for them or for the base asset.

This is one concrete way to discipline remarks like "money will just be energy." Energy, compute, and hardware are inputs; Work Credits denominate future access to verifiable, robot-mediated output. They do not become joules, reserve assets, or monetary anchors by denomination.

12.4 Why Work Credits Remain Service Claims

A claim can be honest and still not be money. Scarcity, transferability, public verification, censorship resistance, and recurring demand are together enough to make a Work Credit a high-quality commodity or forward service claim; none of them, singly or jointly, is enough to make it savings. The reason is physical rather than rhetorical: the credit stays exposed to workload basis, location, hardware, SLA, expiry, delivery, and common-cause risk, and a promise of future service may carry duration. It is a warehouse receipt with a date on it, and a savings instrument has to survive the day the warehouse is the problem. DVC and the evidence artifacts determine whether the claim is fulfillable; the Native Monetary Buyer Map determines whether the separate base asset has a holder constituency capable of bearing loss. Two different questions, two different instruments, and neither answer transfers.

12.5 Monetary Role

Work Credits sit at the junction between **base capacity** and **service contracting**:

- For operators, they are **revenue in kind**: miners/provers/routers earn credits by contributing triad capacity.
- For users, they can be **pre-paid capacity**: hold credits to secure future access to specified triad services, or trade expected scarcity.

Their service value depends on:

- **Demand for specific workloads:** Work Credits tied to high-value workloads (e.g., compliance proofs, LLM inference) may command higher premia.
- **Governance & telemetry honesty:** If VerifyPrice and decentralization metrics are falsified or gamed, the link between credits and real work weakens.
- **Regime pressure and deliverability:** Differential demand may rise as substitutes weaken, while DVC may fall; the Pressure–Capacity Corridor prevents a monotonic crisis claim.

In portfolio terms, Work Credits are typed service or capacity exposure. LP and staking positions are the equity-like or derivative layer; the base asset is evaluated separately.

12.6 Issuance: Tying Credits to Real Capacity

There are many possible issuance schemes; what matters here is not choosing a particular curve, but enforcing two principles:

1. **Issuance is legible.**
2. **Issuance is constrained by real capacity.**

One extreme is the Bitcoin model: a fixed schedule, regardless of demand, with the understanding that price will equilibrate. The other extreme is a pure capacity-linked model, and it is the one the reference design takes: a Work Credit is a capacity reservation, sold when committed capacity is reserved and retired when the service is delivered, inside an issuance envelope set by stress-adjusted DVC.

In such a system, adding a new proving cluster, plant, or corridor is not just a marketing slide; it is an event that expands the envelope of Work Credits the system can credibly support. The converse is also true: if plant retires or corridors die and are not replaced, the reservations against them must move to substitute capacity under their terms or default, and any sales that continued on autopilot would show up as VerifyPrice drift, SLA breaches, and deteriorating energy metrics. The governance layer's job is not to guarantee any particular price, but to keep issuance and capacity in rough proportion and to make any departures visible.

From an allocator's perspective this yields a familiar pattern: Work Credits look like prepaid claims on the output of a portfolio of infrastructure (proving farms, AI chains, privacy corridors, and plants) rather than a purely arbitrary balance sheet. The difference from conven-

tional "infra tokens" is the insistence on receipts and KPIs: if claimed capacity and observable behavior diverge, the discrepancy is not a rumor; it is a datapoint.

12.6.1 The Supply Balance Equation

Here is the question that kills most "energy-backed" tokens in the cradle: **why does demand not simply expand supply, neutralizing scarcity?** The mechanism sounds self-defeating. If Work Credits are sold whenever a customer wants service, then every new customer who arrives creates new supply to serve them. The token is busy, the network grows, and the scarcity — the very thing that was supposed to make holding worthwhile — evaporates into its own success. Demand, on this reading, is not a source of scarcity but an exhaust valve for it.

That is not a bug in the design. It is the design, correctly read. Work Credits are sold against reserved capacity and retired on delivery, so supply growing with demand is not leakage; it is the claim staying honest. A service claim whose supply could not grow to meet service demand would be a shortage instrument, not a savings instrument — and neither of those is a reserve asset. The interesting quantity is therefore not gross issuance but what remains after the service is delivered, and that is a balance, not a schedule.

The answer lies in the **net supply dynamics**. It can be written as a balance equation:

$$\Delta\text{Outstanding WC} = \text{Issuance} - \text{Redemption/Retirement} \pm \text{Governance}$$

Lost or forgotten wallets do not leave the outstanding count — the units still exist on-chain — so they are not a term in this identity; the *effective* (circulating) supply additionally nets them out.

Term	Definition	Typical Magnitude
Issuance	New WC sold against committed, reserved capacity, subject to issuance envelope	Bounded by schedule or capacity ceiling
Redemption/retirement	WC destroyed when the specified service is delivered or the claim expires	Set by contract terms
Lost	WC in lost/forgotten wallets — retained in *outstanding*, excluded from *effective* supply	Modelling assumption of order 1% of supply p
Governance	Adjustments via protocol upgrades (rare, requires supermajority)	Near zero in steady state

For Work Credit supply to remain aligned with deliverable service, the design must

ensure:

$$\text{Outstanding claims} \leq \text{haircut stress-adjusted DVC over the delivery horizon.}$$

Read that as a solvency test, not a scarcity test. It is the warehouse being asked to show that every receipt outstanding corresponds to something physically present and deliverable. It places no ceiling on how many receipts a busy system writes; it places a ceiling on how many can be outstanding at once, and the ceiling moves only when the physical world moves. That distinction carries the whole chapter:

- **During growth:** Work Credit issuance may grow with independently verified DVC and contracted demand.
- **At maturity:** Claims retire through delivery, expiry, or cancellation under published terms. Scarcity can affect service price without making the claim a reserve asset.

What prevents runaway issuance? The same four doors, in the order an operator would try them:

1. **DVC envelope:** Total WC outstanding for each workload, tier, and horizon is capped by haircut stress-adjusted DVC.
2. **Capacity evidence:** The envelope expands only when FERs, FCRs, hardware profiles, and scenario flow support additional deliverable service.
3. **Contract retirement:** Delivery, expiry, cancellation, and default rules remove or impair claims explicitly.
4. **Governance friction:** Changes to service-claim rules require supermajority and are visible in dashboards before activation.

Telemetry that detects supply risk:

- **Inflation rate:** Issuance / circulating supply. Should decline over time.
- **Retirement rate:** retirements / issued. Should track delivered service and expiry terms.
- **Net supply change:** ΔSupply per epoch. Positive during growth, flat or negative at maturity.
- **Issuance vs. capacity:** If issuance grows faster than verified capacity, this signals potential dilution.

If these metrics drift outside healthy ranges, it becomes visible in dashboards—holders and operators can respond before the SoV thesis is undermined. A system that sells against a warehouse cannot hide a shrinking warehouse; the drift shows up in the issuance numbers before anyone has to allege anything.

12.7 Energy-Priced, Not Energy-Pegged

In popular discourse one often hears that "money will just be energy," especially in the context of robotics and AI. As shorthand, this is attractive: robots and data centers run on electricity, so why not quote everything in kWh and be done with it? The problem is that not all kilowatt-hours are created equal. Time of day, grid node, reliability, carbon intensity, and siting constraints all affect their economic and political meaning. A winter-peaking kWh on a stressed urban grid is not the same as a curtailed hydro kWh in a remote valley. Pretending otherwise smuggles a lot of hidden politics and risk into the unit of account.

The discipline in this thesis is to admit that complexity and route around it with **receipts**.

Energy is measured and wrapped into Facility Energy Receipts (FERs), which record, for each facility and time window:

- kWh in,
- kWh delivered to IT,
- heat reused,
- PUE/ERE/WUE,
- water use,
- carbon intensity,
- and outages/curtailments.

On top of that, the thesis measures **verified work per kWh**:

- η_{vFLOP} for FLOPs,
- proofs-per-kWh for proof workloads,
- swaps-per-kWh for settlement.

VerifyPrice then tells us how much it costs, in time and money, for an independent verifier to check that a given amount of work was done.

The result is an implicit conversion path:

> **energy → FERs → verified work (proofs/FLOPs/swaps) → receipts → Work Credits**

The design never pretends that *one Work Credit is one kilowatt-hour*; instead it makes it easy for anyone to estimate, at any given time, **how many kilowatt-hours of which quality and where in the world** sit behind a portfolio of Work Credits, via the receipts.

Pricing in energy then becomes an **inference problem for markets**:

- Work Credits are implicitly energy-priced because their production and redemption depend on energy-intensive workloads whose cost curves are public.
- FERs and VerifyPrice make those curves legible without forcing a brittle kWh peg.

What distinguishes Work Credits from naive "energy tokens" is precisely this separation:

- The unit of account is denominated in **work**, not in kWh.
- Energy enters through FERs, η_{vFLOP}, and VerifyPrice, not through a one-dimensional peg.

Markets, regulators, and builders can look at those receipts and say, with some confidence, "a Work Credit currently corresponds to about this much capacity, with this energy and carbon profile, under these SLAs." That is enough to make the asset priceable and analyzable without forcing it into a crude energy standard.

12.8 Why Work Credits Are Typed Service Claims, Not Presumed Money

It is equally easy to make the opposite error: to observe that Work Credits are scarce, transferable, and tied to necessary workloads, then promote them from service claims into money by vocabulary. This thesis does not make that move. A Work Credit is a typed claim on capacity or service whose quality depends on workload, location, hardware, SLA, delivery window, and the stressed path that can actually fulfill it. That case was made at Section 12.4; this section asks the harder question it left open. What would have to be *true* — not said, not branded, but observable — for a Work Credit to cross from service claim into money?

Four things, each visible in telemetry before it is visible in price:

1. **Issuance is constrained by stress-deliverable capacity, not nameplate capacity.**
   FERs, FCRs, and Delivered Verified Capacity bound claims by the service that survives scenario-specific binding constraints. Gross installed hardware, energy procurement, or benign-state output is insufficient.
2. **Demand can be budgeted without being monetary.**
   Proofs, inference, settlement, and related workloads may sit on recurring OPEX lines. That supports a service market. It does not establish reserve demand, long holding periods, self-custody, loss-bearing capacity, or countercyclical accumulation.
3. **Verification makes the claim legible, not monetary.**
   A PIDL receipt can make fulfillment cheap to check and FCR/FER artifacts can make provenance auditable. Evidence reduces information asymmetry; it does not erase delivery, basis, counterparty, or temporal risk.
4. **The claim must not absorb project duration by implication.**
   A prepaid capacity claim, a claim on future service, and a project note are different instruments. If a Work Credit promises par, redemption, a coupon, or emergency support against construction failure, it has become duration-bearing credit and must be labeled and underwritten as such.

So what would the crossing actually look like? Not a marketing decision and not a milestone on a roadmap, but a change in *holding behavior* that the dashboards can see: holding periods lengthening rather than claims being redeemed at term, balances retained through drawdowns instead of recycled into service, and — the tell that would settle it — buyers acquiring the asset *while the service it claims is cheap and available*, which is the signature of savings

rather than procurement. None of that is a design property. It is holder behavior, and it is exactly what the Native Monetary Buyer Map watches for on the base asset (Section 26.8.8). Until those series move, Work Credits belong beside capacity vouchers and service contracts, not beside the base monetary candidate. They may trade at a premium for scarce delivery rights and may serve as collateral where their terms permit. Those are properties of typed service claims. Monetary treatment remains a separate empirical question answered, if at all, by the holder-side mechanism and buyer map (Section 13.3.11).

Chapter 13

Why Hold the Base Asset?

A holder who wants verified service in a bad state can hold one of three things: a Work Credit, which is a typed claim on reserved service owed by a named provider (Chapter 12); Bitcoin, to buy what is needed when the need arrives; or the base asset. The first two are well understood. The case for the third has to be made against them, and it is made here in the order a skeptical reader would demand: the strongest alternative first, then the channels through which service value reaches the asset, then the reasons that value might be monetary, then the instruments that would show it is not.

13.1 Why Hold Rather Than Convert on the Day?

Where payment requires the base asset, a Bitcoin balance remains dependent on acquisition until that acquisition has settled. A balance held in advance removes that dependency. That is the whole of the advantage, and it exists only in a band: where conversion is unavailable, too slow, or incompatible with the transaction's required privacy and control conditions, while native settlement and the complete service path remain usable. Outside the band, the holder of Bitcoin buys what is needed when it is needed and carries none of the base asset's risk in the meantime.

Two different advantages are easy to run together. The *stack's* advantage is that it can deliver what the best institutional substitute cannot. *Preholding's* advantage is that a holder can obtain that delivery when the best Bitcoin-funded strategy cannot. The first does not establish the second. Section 30.6.1 explains why a Work Credit hedges the wrong state and why Bitcoin does not itself produce a proof; neither argument defeats Bitcoin-funded procurement.

The two paths are

$$\text{held native balance} \;\rightarrow\; \text{native payment} \;\rightarrow\; \text{service},$$
$$\text{Bitcoin balance} \;\rightarrow\; \text{acquisition or other procurement} \;\rightarrow\; \text{service},$$

and the mechanism exists only where the extra dependency in the second path makes a material difference, not because the second path has one more arrow. For outright timing failure

the condition is

$$T_{\text{native}}(s, W) \;\leq\; \tau \;<\; T^{*}_{\text{BTC}}(s, W),$$

where s is the stress state; W carries the workload, assurance tier, privacy requirements, and delivery terms; τ is its deadline; T^{*}_{BTC} is the completion time of the best feasible Bitcoin-funded strategy; and both strategies start from the same purchasing-power budget. The broader case includes differences in cost and in the probability of failing, not only outright inability to transact, but every such difference must survive the cost and risk of holding the base asset beforehand.

Where the band might be. The states in Table 13.1 are hypotheses registered in advance, not examples of the mechanism working.

Stress state	**What fails for buying on the day**	**What must still work for a held balance to help**
Selective acquisition restrictions	Accessible venues withdraw service, demand credentials the buyer cannot supply, or impose disclosure the transaction's terms forbid	The holder already controls a spendable native balance; native admission does not reproduce the same restriction; an acceptable provider remains reachable
Conversion-liquidity interruption	Executable inventory disappears, withdrawals stall, or the quantity needed cannot be bought within the workload's deadline and budget	Native settlement works; enough providers accept the asset; they can keep operating without the interrupted conversion channel
Asymmetric settlement disruption	The source-asset or cross-system leg cannot reach the finality the deadline needs	Native settlement and the complete service path remain live
Deadline-bound delegation	A delegation can still procure and verify work, but an extra acquisition round trip exceeds its authorized window	A balance authorized in advance, and revocable by its human principal, can pay inside the window without widening the principal's permissions

Table 13.1: Pre-registered stress states for the preholding mechanism. Hypotheses to be tested, not observed cases.

The last row gives the delegation buyer of Chapter 7 a concrete role, with a limit: it is a reason for prefunded transaction balances, not for monetary reserves.

What the balance buys there, and what it does not. It buys payment capability that is immediately usable inside the service economy that survives. It does not buy a reserved quantity of compute, a fixed service price, a queue position, a provider's continued existence, or protection against the asset losing purchasing power. Providers promise no capacity, price, or priority because a customer holds the asset. Those rights are sold through specified service contracts (Chapter 12), and the contract, not possession of the payment asset, defines the obli-

gation. That is why the base asset carries no claim on historical or future work, and why it is compatible with a provider accepting payment at the moment of procurement. Acceptance by providers is part of the service path, not a separate guarantee: even universal acceptance leaves open why the holder should not acquire the accepted asset immediately before paying. The condition is not acceptance by every surviving provider but enough independently supplied, reachable capacity accepting native payment to fulfil the workload on its terms; a provider that stays powered and refuses the asset contributes nothing to it.

The counterfactual is allowed to win. The comparison is against the best Bitcoin-funded strategy actually feasible in the state, not a weak one such as opening an exchange account during the crisis. That strategy may use prepared wallets, existing non-custodial routes such as the corridor of Section 23.4, direct payment to providers, funded payment channels, independent conversion providers, and diversified service reservations. Above all, it may let the provider do the conversion: a protocol that requires native fees does not thereby require the customer to hold the fee asset. Two results would damage the mechanism, and both must be visible if they occur. First, a provider or router may accept Bitcoin and pay the fee from its own native inventory. The system then needs native inventory somewhere, but not a durable reserve held by the end user. That is the bypass that Conditions 1 and 5 of the Value Capture Lemma name, seen from the holder's side, so this test is also the non-bypassability test. Second, providers that depend on selling native receipts to pay for power, hosting, and hardware may stop operating when the same acquisition market fails, so a disruption that first favours holders can then disable the service they hold for. The test therefore runs over the delivery horizon, not to the first successful payment, and the Pressure–Capacity Corridor (Section 8.2) still governs: no balance insures access through a payment path whose suppliers stop during the same event.

The access experiment. The holder service is tested before it is claimed. Equal-budget strategies are fixed before the disruption: a held native balance; Bitcoin with acquisition on the day; Bitcoin with provider-side conversion; an existing reserve asset plus diversified service reservations; and direct procurement where it exists. Each faces the same service request under the same adversarial scenario. The published readings are completed work, completion time, privacy and control failures, total procurement cost, and the losses borne during the holding interval, with balances held in anticipation of stress reported apart from balances acquired once the stress was obvious. A pass requires more than a native transaction succeeding. It requires a material advantage over the strongest alternative that was actually feasible. The failure condition is stated now: **if Bitcoin-funded procurement reproduces the same service, within the same constraints, without a material disadvantage, the preholding mechanism is absent for that workload and state.** The experiment is a Phase I deliverable of Chapter 32

(Table 32.1), run first as an analog on an existing network whose fees must be paid in its own asset. It is not a counted red line. It is the evidence Red Line 16 stands on: the experiment asks whether holding changes what can be done, and the lending and basis series of Red Line 16 ask whether a mature market prices the holder service (Section 30.6). Neither substitutes for the other.

A pass is not yet a reserve. Even a positive result leaves the second question open: does the benefit support persistent, loss-bearing reserve demand, or only an operating buffer? Agent wallets, gas inventories, and precautionary balances do not become a monetary anchor by being added up; that is Premise 7, and it stands (Section 1.2). Whether independent principals keep balances beyond their transaction windows, and bear the risk of them through stress, is what the Native Monetary Buyer Map (Section 26.8.8) and Red Line 16 read.

> The base asset is not a reservation on the machines. It is a balance that may remain spendable when acquiring that balance would no longer be feasible. Whether that difference survives the best Bitcoin-funded alternative, and supports anything beyond operating inventory, is the monetary experiment.

13.2 Economic Linkage: How Triad Demand Becomes Asset Value

One question separates a *systems manifesto* from a *monetary thesis*: **why does demand for Privacy, Proofs, and Compute raise the value of *holding* the asset, rather than merely rewarding *consuming* a service?**

The familiar answer is a chain that runs from fee demand through burns to scarcity to price. That chain fails (Section 13.3.9) and is not used here in any form. The answer below is in four parts, of which the first is a cash-flow claim the thesis does not describe as monetary, and the remaining three are the monetary content.

1. What exactly is the asset? The **base asset** is the protocol-internal unit in which fees may be required, collateral posted, and settlement denominated. It is a conditional monetary candidate and nothing in this section upgrades it. Work Credits are typed capacity or service claims; FCR, FER and PIDL artifacts are evidence; project notes are duration-bearing credit; LP and staking positions are derivatives or operating claims (Section 8.4). Only the base asset is evaluated as money.

The strong phrasing — that *all* uses of the triad must flow through the asset — overstates the design. Core protocol operations are denominated in the asset by rule. Equivalent computational content is obtainable elsewhere, always, at a price; bypass is a magnitude rather than a switch (Section 13.3.4), and the design constrains what that magnitude is worth rather than

abolishing the alternative.

2. Why is it scarce in a monetary sense? Scarce *capacity* does not imply a scarce *asset*. The design forges the link through scheduled issuance, fee retirement, collateral lockup, and Work Credits sold against committed, verifiable capacity rather than by decree.

Two precisions the earlier text elided. First, base-asset issuance follows the constitutional schedule and is independent of Work Credit dynamics; it is Work Credits that are bounded by delivered verified capacity. Second, whether burns exceed issuance at steady state — whether net supply actually falls as workloads grow — is an empirical outcome the fee-coverage and issuance boards report, not a property the design confers. Issuance discipline is real. Net deflation is a reading.

3. Why does demand for triad capacity raise the value of holding? Four channels, stated in ascending order of monetary content, with the first two explicitly conceded not to be monetary at all.

1. **A fee stream (not monetary).** Every proof, private settlement and verified inference pays a fee in the native unit, and routing rules direct it to retirement, staking, and operator share. What the holder owns through this channel is a competitively priced claim on a cash-flow stream bounded above by differential value Δ and deflating with unit price (Section 13.3.3). A discounted-cash-flow valuation reproduces it exactly. A burn is a buyback and confers no moneyness.
2. **A collateral lockup (not monetary).** Required lockup in the native unit immobilizes float while operators remain solvent and choose to keep operating. Section 13.3.12 bounds how large a fact this can be — large-cap-equity magnitudes, not monetary aggregates — and shows it is a slope effect on price impact rather than a level effect on price, because a lockup places no bid.
3. **Pledgeability without due diligence (monetary).** This is the first channel through which growing triad demand raises the value of *holding* rather than of *receiving*. As verified throughput grows and receipts accumulate, the cost to any counterparty of establishing what a pledged position actually is falls toward zero, and a position nobody must investigate can be accepted no questions asked (Section 13.3.10). The holder is not being paid; the holder is holding something that has become easier for everyone else to accept. Nothing is distributed through this channel, so a valuation of distributions alone omits it. Whether a holding service of this kind is monetary is what Section 13.3.10 and Red Line 17 test.
4. **A unit with something to denominate (monetary).** A growing endogenous service economy gives counterparties a reason to reckon in the unit rather than merely settle in it, and denomination is the stickiest monetary function because it is a coordination device rather than a portfolio choice (Section 13.7.3). This channel is available to a service protocol and is

not available to a bearer asset with no native goods.

A fifth channel, the state-contingent holder-side service flow, is argued separately at Section 13.3.11, because it is a property of the regime rather than of demand growth.

> Channels 1 and 2 pay the holder. Channels 3 and 4 change what holding *is*. Only the second kind is monetary, and only the second kind survives a correctly specified discounted-cash-flow valuation.

A holder of channels 1 and 2 alone is buying "a share of the fee stream from indispensable workloads," a position "closer to equity in a utility than to a collectible." That description is accurate for those channels, and it is the reason they cannot carry the monetary reading: a utility equity earns its cost of capital and carries no monetary premium. Channels 3 and 4 are the reply, and they are a reply about the asset's acceptability and denomination rather than about its cash flows. Neither is sized in this document.

4. What prevents capacity providers from capturing all value? The classic "utility token trap": if operators earn the fees and governance can inflate supply, holders are exit liquidity. The design's answers are routing rules and they are stated as such — fees split between retirement, stakers and operators so no party takes the whole; base-asset issuance constitutionally scheduled; Work Credit sales bounded by reserved capacity, with overselling rejected and visible; core operations denominated natively.

These govern *where the fee goes*, and none of them governs how large the fee can be, which is the question Section 13.3 takes up and answers with a bound rather than a promise. Read this paragraph as a description of plumbing. The monetary argument is in channels 3 and 4 above and in Section 13.3.11, and it does not depend on any number in this paragraph.

Summary: the Value Capture Lemma in plain language

> Demand for triad capacity becomes value *accrual* to the native asset only if five conditions hold at once: all core usage requires the asset, usage retires supply, capacity provision requires staking, issuance is constrained, and capacity cannot be cheaply bypassed. These conditions are necessary and not sufficient, and what they establish is accrual — a cash-flow claim and a collateral lockup — not moneyness. If any one fails, the system may be indispensable infrastructure whose economics accrue to operators or customers.

A related but separate claim. The five conditions are properties of the monetary design and say nothing about the durability of the demand being routed. That workload demand is *structural* rather than speculative is a distinct empirical claim whose failure the thesis treats as its own falsifier: a workload mix that becomes and remains predominantly speculative retires the store-of-value claim through Red Line 5, not through this lemma.

Lemma 13.1 (Value Capture Lemma)**.** *In the reference design, demand for Privacy, Proofs, and Compute* accrues economically *to the native asset — as a claim on the fee stream and a collateral lockup — through five channels, and the design relies on all five holding together:*

1. ***Required fee medium:*** *the asset is required for core fees, and users cannot pay in fiat, stablecoins, or other tokens at equivalent service quality.*
2. ***Supply reduction:*** *a meaningful share of fees is burned or permanently retired.*
3. ***Collateral lockup:*** *operators must stake the asset as collateral to provide services.*
4. ***Issuance discipline:*** *gross issuance follows a schedule fixed ex ante, not governance fiat.*
5. ***Non-bypassability:*** *users cannot obtain equivalent triad capacity through bypass channels (cloud contracts, stablecoin-denominated services, direct fiat payment to operators) that avoid touching the asset.*

Each channel closes a named leak, and in this design nothing else stands in for a channel that fails. The five are this design's route to accrual, not the only route economics allows: a protocol can take fees in another unit and use them to buy and retire the asset, distribute revenue instead of burning it, or take native fees without native collateral, and each of those captures value by a different route with different risks. The lemma is the bridge between "these capacities are indispensable" and "therefore this asset accrues the value of that indispensability." Accrual is a precondition the monetary claim builds on, not the monetary claim itself (Section 13.3.9, Section 13.3.11).

Argument. Each condition closes a leak that its negation names. If fees can be paid in another unit at equal quality (¬1), the marginal user pays in that unit and the fee stream accrues to whoever accepts it — unless the protocol converts those fees into the asset and retires it, which is a buyback funded from outside the native unit and a different design. If no share of fees is retired (¬2), the stream accrues to operators and to the stakers who work for them as income, and a holder who neither operates nor stakes owns a claim on nothing — unless the design distributes revenue pro rata instead, which creates a financial claim by a different route and carries the classification exposure of Section 27.2. If provision requires no lockup (¬3), nothing immobilizes float, and the asset's price can fall to zero without any operator's service being interrupted. If issuance is discretionary (¬4), any accrual can be diluted at will and a counterparty must hold a private view of the issuer, which is the information-sensitivity failure of Section 13.3.10. If equivalent capacity is obtainable without the asset (¬5), Condition 1 is satisfied only nominally and the fee is bounded by the bypass cost, which Section 13.3 shows can be zero; a provider that accepts Bitcoin and pays the fee from its own inventory is this bypass seen from the holder's side, and the access experiment of Section 13.1 measures it. The argument shows that each condition is necessary *to this design's route*, not to value capture in general. Sufficiency is not claimed either: all five can hold while the fee f is zero (Section 13.3) or while accrual occurs and no monetary premium follows (Section 13.3.9). □

Which conditions the monetary channels actually use. A reconciliation that the earlier text could not state, because it had only one candidate mechanism. Of the five conditions, four govern *accrual*, which is a discounted-cash-flow quantity. **Issuance discipline is the exception**, and Section 13.3.10 explains why: an asset whose supply is discretionary is information-sensitive on supply, because a counterparty must maintain a private view of what the issuing body will do, and an asset requiring a private view cannot be accepted no questions asked. Condition 4 is therefore doing monetary work that Conditions 1, 2, 3 and 5 are not. That is the same conclusion Section 13.3.11 reaches from the other direction, and the two now agree for a stated reason rather than by coincidence.

What the five conditions do not address. Every one of them governs *where* fees go. None governs *how large* they can be, and a routing rule applied to a fee of nothing captures nothing. Section 13.3 takes up the level question, finds that the standard objection to taxing a contestable market is wrong for a reason worth stating, and then finds that the correct answer supports a narrower conclusion than this Part had been drawing. The conditions themselves are unchanged.

The lemma has a frequently-missed converse. It states when native demand accrues *economically* to the asset. It does not state that movements in the asset's *observed market price* are caused by that demand. Those are different machines, and conflating them is the most common way a monetary thesis fools its own author.

Corollary 13.1 (Market-Price Non-Equivalence)**.** *External financial wrappers can create price demand without exercising privacy settlement, purchasing proofs, consuming verified compute, paying native fees, locking collateral, or producing burns. Conversely, wrapper redemptions, leverage unwinds, dealer hedging, or broad risk-off flows can depress the asset's price while native monetary usage improves. Therefore*

$$\textit{price appreciation} \not\Rightarrow \textit{monetary adoption}$$

and

$$\textit{monetary adoption} \not\Rightarrow \textit{immediate price appreciation}.$$

*Price is an output to explain, not a protocol KPI to worship. The monetary thesis must be evaluated through native fees, burns, collateral, settlement, receipt volume, verification health, reachability, and non-custodial use. Market-price interpretation requires a separate financial-flow telemetry layer (**VerifyFlow**, Section 26.8).*

Argument. The market price is set by the marginal trade in a market that includes every instrument in Table 8.4. Wrappers below the base row exist and trade without exercising any native function, so a price move can be produced by wrapper flow alone: price appreciation is therefore not sufficient for monetary adoption. Conversely, the same instruments can be unwound while native series improve, so adoption is not sufficient for appreciation. The corol-

lary follows from the existence of the wrapper rows, not from any assumption about their size; Appendix H quantifies the size. □

13.2.1 Value Capture vs. Price Capture

It is worth separating the two phenomena explicitly, because the thesis is routinely "confirmed" by the wrong one.

Value capture is the existing loop: triad usage → native fees → burns, collateral lockup, and operator demand → scarcity and *cash-flow accrual* to the asset. It is governed by Lemma 13.1. What it produces is a claim on a fee stream plus a collateral lockup; whether the asset also carries monetary premium is a separate question, answered separately in Section 13.3.11.

Price capture is demand for exposure to the asset's price, arriving through spot ETFs, exchange custody, treasury companies, margin products, options, futures, swaps, leveraged and inverse ETPs, passive indices, and systematic trading rules. Price capture can occur with no protocol use whatsoever.

Native value capture	Price capture	Interpretation
Strong	Strong	Service-value accrual with favorable market realization
Strong	Weak/negative	Protocol improving while wrappers or risk markets sell
Weak	Strong	Financialized speculation or wrapper-led adoption
Weak	Weak	Weak accrual and weak price; silent on the monetary mechanisms

Table 13.2: Native value capture and observed price appreciation are neither equivalent nor mutually necessary. No cell is monetary evidence: that is a separate axis, read on Red Lines 16–18.

The diagonal is what makes the thesis falsifiable. A thesis that treats every price rise as confirmation and every price fall as noise cannot be wrong, and therefore cannot be right either.

13.3 Fee Incidence and the Level of the Fee

Every condition of the Value Capture Lemma answers one question: *where does the fee go?* Required medium, supply reduction, collateral lockup, issuance discipline, and non-bypassability are all routing rules. None of them answers the question a valuation actually needs answered, which is *how large the fee can be*. A protocol that routes the whole of a fee of nothing into a burn address has satisfied five conditions and captured zero.

The gap has a sharp form, best stated as the objection it invites:

> *"You have deliberately engineered the most contestable market you could. Workloads are canonical and interchangeable, verification is cheap enough that no trust is required, admission is open, and there are no vendor chokepoints. A contestable market competes economic profit to zero. Zero profit means no surplus, and no surplus means nothing to burn."*

The objection is wrong. It is worth showing exactly why, because the correct answer rescues considerably less than it first appears to, and what survives is not the claim this Part started with.

13.3.1 Incidence: A Fee Is a Wedge on Turnover, Not a Claim on Profit

The objection conflates two different bases. A protocol fee is levied on *transactions*, not on *residual profit*. It is a tax wedge, and tax wedges do not require anybody to be earning rents.

Take the market the thesis is trying to build: many providers, free entry, roughly constant returns to scale in the relevant range, a canonical workload so that units are fungible. Long-run supply is then approximately horizontal at marginal cost c (equivalently the protocol's all-in unit cost c_p of Section 13.3.3; write $c \equiv c_p$). Impose a per-unit protocol fee f on that market and the supply curve the buyer faces shifts to $c + f$. The new equilibrium has buyers paying $c+f$, providers receiving c, and providers earning exactly the zero economic profit they earned before. The protocol collects $f \cdot Q$. Nothing about that revenue came out of operator margin, because there was no operator margin to take it from; in the long-run competitive case the entire statutory burden lands on the buyer.

> Operator margin compression is irrelevant to fee revenue. The fee is a wedge on turnover. Turnover is what a contestable market maximizes.

This is standard incidence analysis, and it means the objection proves too much: if contestability destroyed fee bases, no commoditized market could ever be taxed, and every one of them is.

Incidence explains who *bears* a fee. It does not explain why a positive fee is *sustainable*, and the two should not be run together. Exchanges, clearing houses, and toll roads are lucrative on top of viciously competitive flows for a further reason: the flow is competitive but the fee layer is not. There is one place to clear the contract, one road across the water, one book where the trade is legally final — so the competition among users of the layer never becomes competition to *be* the layer. That inescapability is a property those venues possess and a protocol must earn; it is the subject of Section 13.3.3 and Section 13.3.7, and the fork channel is precisely the case where it is not earned.

The short run is different, and the difference matters. With installed capacity fixed, supply is not horizontal. Deployed hardware has already sunk its capital, so in the short run the supply curve slopes upward and part of the fee *is* absorbed as a reduction in quasi-rent to existing machines rather than passed to buyers. That is not a problem for this quarter's fee revenue; it is a problem for the next round of capacity. Quasi-rent is what funds replacement, and an operator whose realized return on deployed hardware is being clipped by a fee re-prices the entry decision accordingly. The long-run pass-through result therefore describes the destination, not the transition, and the transition is where entry and capacity growth are decided. A fee schedule set as though incidence were instantaneously long-run will suppress the capacity it is trying to tax.

13.3.2 The Base, and Why It Deflates

Fee revenue is not f. It is

$$\text{Rev} = f \cdot Q(c_p + f),$$

and both terms are endogenous. Two consequences follow, and the second is the more serious one.

Quantity responds to the wedge. The tollbooth analogy holds only where there is no parallel free road. Raising f raises the delivered price, and at the margin some volume declines, defers, batches, or leaves. What the protocol can extract is bounded not by what buyers would pay for the service but by what they would pay *through this channel rather than the next-best one*, which is the subject of Section 13.3.3.

An ad valorem fee on a deflating unit price is a shrinking real toll. For this paragraph the protocol's fee is modeled as *ad valorem*: a share of unit price rather than the per-unit wedge f used in the incidence analysis above. Under a per-unit fee only the volume term below moves, and the deflation mechanism this paragraph isolates does not arise. Section 11.1 argues that verified compute stays *scarce* even as raw compute deflates. Grant that argument entirely; it does not touch this one. Even if the verification premium is fully durable, a fee expressed as a percentage of a unit price that falls secularly collects fewer real resources per unit of service delivered every year. Real fee revenue then grows only if

$$\underbrace{\text{growth in verified units delivered}}_{\text{volume}} > \underbrace{\text{rate of decline in real unit price}}_{\text{deflation}}.$$

That inequality is an empirical bet, not a theorem, and the thesis is making it. It may well be a good bet — the historical pattern in compute is that unit-cost decline is met by more than proportional volume expansion — and a bet must say what would settle it against the thesis.

Condition B of Red Line 14 (Section 30.6) does: deflation-adjusted gross native fee turnover that fails to grow while physical throughput grows is this bet losing, in public.

The bet has a macro frame. The volume-over-deflation bet is usually argued from technology curves: unit costs fall, demand expands more than proportionally, revenue grows. That is the technological channel, and it is the one Section 11.1 treats. There is a second channel with the same signature but a different driver: a financial clearing event that transfers installed capacity from weakly financed owners to strongly financed users at lower capital cost, collapsing the effective price of capacity without any change in the technology [Perez 2002]. The two channels are observationally distinct. The technological channel shows up as gradual unit-cost decline with stable utilization, the financial one as a step-change in the price of *installed* capacity with utilization resetting. They have the same effect on the inequality above: both push the deflation term up. What the frame adds is timing. The technological channel is smooth; the financial channel is not, and a fee base built on the expectation of smooth deflation being outrun by volume will experience the financial channel as a shock to exactly the term it is betting on. The instrument is unchanged, since Condition B reads on the same series either way. But a reader tracking it should know that the two channels are live simultaneously, and that the financial one can move a quarter's reading by more than a year of Moore's law.

13.3.3 The Bound on the Sustainable Fee

Incidence tells us the fee has a base. It does not tell us how high the fee can be set. That ceiling is not a function of operator margin either; it is a function of what the protocol channel is *differentially worth*.

Differential Value Δ and the Sustainable Fee Bound

Let c_p be the protocol's all-in cost of delivering canonical workload W at tier T, and c_b the all-in cost of the best bypass channel delivering the same computational content without protocol-grade verification. Let Δ be the marginal buyer's willingness to pay for the protocol's differential properties: cryptographic checkability, neutrality, credible non-discretion, and non-custodial exit.

The buyer routes through the protocol only if $c_p + f - \Delta \leq c_b$, so the sustainable fee satisfies

$$\boxed{f \;\leq\; \Delta - (c_p - c_b)}$$

Here Δ is a scalar wedge in price units, not the difference operator used elsewhere in this Part.

The first term is the one a protocol designer wants to think about. The second is the one that decides the outcome, and it is **structurally positive**: verified compute costs *more* to produce

than unverified compute, because proving overhead and redundancy are real resources spent on top of the underlying work. So the protocol begins every negotiation with a cost handicap and must recover it out of Δ before it collects a single unit of fee.

The document's own tier multipliers are cost pass-through, not evidence of willingness to pay. Section 22.3.2 prices assurance on its own axis: Replicated at $1.0\times$, Replicated-plus-audit at $1.4\times$, Probabilistic (M2) at $1.6\times$, Succinct (M1) at $2.5\times$, with a latency multiplier of $1.0/1.3/1.6$ multiplied in. Read those against what each tier costs to supply. Replicated-plus-audit is three independent re-executions and a randomized audit, so its multiplier is set against roughly three times the production cost; the M1 multiplier is set against a proving overhead that Appendix E.6 publishes as a quarterly series and does not yet know. The multipliers are rationale-set to cover redundancy or proving overhead plus a margin, and the document nowhere establishes that they do. That is a question about c_p and the margin above it, not about demand, and nothing in the document establishes that any buyer's Δ for succinct assurance is $2.5\times$ its Δ for a committee's word. If it is not, headroom *narrows* as assurance rises, and at the top of the axis $\Delta - (c_p - c_b)$ can be negative — a tier that is technically superior, honestly priced, and unsellable. That is a testable prediction against the thesis and it is now measured (Section 30.6, Red Line 14), because it should not have been left as an inference from a table caption.

13.3.4 Bypass Is a Magnitude, Not a Binary

Condition 5 of the Value Capture Lemma is written as an existence condition — users *cannot* obtain equivalent capacity through a bypass channel — and Red Line 6 is written as a direction condition: usage grows while native fees do not. Neither is wrong, and neither is changed here. But both are read, in practice, as though bypass were a switch.

It is not. Bypass is never impossible. It is priced. There is always a channel that delivers the same computational content without protocol-grade verification, and the only question is what the buyer gives up by using it — which is exactly Δ. Condition 5 should therefore be *read and measured* as a magnitude claim: not "no bypass exists," which is false everywhere, but "Δ is large enough, for enough of the volume, to support a fee worth burning." The condition stands as stated. Its *interpretation* is what this section tightens, and its measurement is what Red Line 14 supplies.

This is why the chain-strength rating for non-bypassability is Medium rather than Strong (Table 8.1). The condition is well specified as a binary and the binary is well enforced. The economics requires a magnitude, and the magnitude is nowhere estimated. A rating of Strong would rate the enforcement machinery, not the quantity the argument depends on.

13.3.5 What Determines Δ, and Why It Is Probably Thin

To first order, the differential value of a proof is the loss it prevents, times the probability that loss survives the alternative:

$$\Delta \approx \varrho\,\mathcal{L},$$

where $\mathcal{L}$ is the loss from undetected misexecution or denial of service and ϱ is the residual probability of that loss under the best bypass channel. This operationalization *lower-bounds* Δ. It prices only loss-prevention, while the full definition in Section 13.3.3 covers four properties: cryptographic checkability, neutrality, credible non-discretion, and non-custodial exit. Neutrality and non-discretion therefore enter as additional terms the $\varrho\,\mathcal{L}$ form omits. (ϱ and $\mathcal{L}$ are used rather than the conventional π and L, which denote the inflation rate in Section 3.1 and the leverage multiple throughout VerifyFlow and Appendix H.)

Δ is a distribution, not a number. Before evaluating either term it is worth fixing what object Δ is. Every buyer has their own ϱ and their own $\mathcal{L}$, so Δ is a distribution across the buyer population. The fee bound of Section 13.3.3 binds at the *marginal* buyer — the last one the protocol wants to retain at the posted fee — not at the median one. These come apart, and the difference is the whole commercial question. A thin median is entirely consistent with a lucrative fee, provided the upper tail is deep enough to be worth serving. That tail is made of counterparties nobody will indemnify because they are sanctioned or pseudonymous, flows crossing jurisdictions where no single forum will hear the claim, regulated model-risk and audit functions that must evidence correctness to a supervisor rather than merely believe it. For those buyers ϱ is not small, because the indemnitor substitute is unavailable rather than merely inferior. The protocol's commercial problem is therefore not "is the median workload willing to pay," which it plainly is not, but "is the high-Δ tail large enough, and reachable enough, to carry a fee." That is a sizing question the thesis has not answered, and Red Line 14 measures it capacity-weighted across the served mix rather than assuming it.

And at the median the answer is unflattering. Written as $\varrho\,\mathcal{L}$, the critical term is ϱ, and here the thesis has been flattering itself. The bypass channel is not "unverified." It is *reputationally and contractually* verified. A hyperscaler's alternative to a cryptographic proof is its balance sheet, its SOC 2 and ISO audit reports, its published incident history, and an indemnity clause enforceable in a court that will actually hear the case.

> A creditworthy, suable indemnitor is an excellent substitute for a cryptographic proof.

For an ordinary commercial workload in a benign state, the buyer's exposure to undetected misexecution is not $\mathcal{L}$; it is $\mathcal{L}$ net of a contractual recovery from a solvent counterparty under a legal system that functions. That makes ϱ small, and therefore Δ thin over the bulk of the

distribution. Thin Δ against a structurally positive $(c_p - c_b)$ is the strongest form of the objection this section opened with, and it is not answered by incidence analysis. It has to be stated plainly, because a critic will otherwise state it first and more crudely: **for the median commercial workload in a working legal order, cryptographic verification is a premium product competing against a cheaper substitute that is good enough.** What that sentence does *not* say — and what a fair reading of it must preserve — is that the median is the marginal buyer. It is not, and the fee is set against the tail.

13.3.6 The Tail Is Bimodal, and the Two Modes Are Mutually Corrosive

The fee is not set at the median. Section 13.3.5 put it at the marginal buyer in the upper tail — counterparties nobody will indemnify, flows no forum will hear, supervised functions that must evidence correctness rather than merely believe it — and then asked whether that tail is large enough to carry a fee.

That is the wrong question, because the tail is not one population. It is two. They want opposite things, and the thesis has been quietly adding them together.

The Two High-Δ Segments

Segment S (supervised). Regulated model risk, audit, clinical and pharmaceutical attestation, regulated AI deployment, supply-chain provenance under statutory duty. ϱ is high because the loss is regulatory rather than financial: "we relied on the vendor" is not an answer a supervisor accepts, and no indemnity cures a sanction. What Segment S requires is *evidence*: scoped disclosure, viewing keys, named counterparties, jurisdictional clarity, retention.

Segment X (unsuable). Cross-jurisdictional, pseudonymous, sanctioned, capital-controlled, or adverse to the forum itself. ϱ is high because the indemnitor substitute is *unavailable* rather than inferior — there is no solvent counterparty a court will make pay. What Segment X requires is *absence*: large anonymity sets, no disclosure surface, no perimeter identification, nothing that constitutes a subpoena target.

Both have genuinely high Δ. Each one's requirement is the other's disqualifier.

The corrosion is not about the rail. It is about the set. An anonymity set is a room. Segment S cannot afford to be seen entering a room that contains Segment X; Segment X is worth nothing in an empty room. Both want the same room, and each wants different people in it.

The obvious objection is that the thesis has already solved this: lawful privacy is default-private with optional disclosure, so Segment S turns viewing keys on and Segment X leaves them off, and both use the same protocol. That answer is correct at the level of the *rail* and it does not reach the actual problem, which is one level down.

Privacy is a network good. What Segment X buys is not a cryptographic property of its own transaction but the size and composition of the crowd it is indistinguishable within, and that crowd is shared. So is the crowd's legal character. The two segments do not merely coexist on a rail; they occupy one another's anonymity set, and each one's presence changes what the set is worth to the other:

- **S needs X to be absent.** A compliance function cannot knowingly transact into an anonymity set containing designated parties, because the commingling *is* the sanctions exposure. The disclosure it can make about itself does not discharge a duty about who else is in the pool.
- **X needs S to be present.** Section 30.6, Red Line 7, already states the mechanism: privacy that only outlaws use is a shrinking set. It is precisely the ordinary, legitimate, high-volume users who make blending plausible, and Segment S is the largest available source of them.

The dependency therefore runs one way and the repulsion runs the other. Segment X wants Segment S in the pool; Segment S wants Segment X out of it. That is not a stable composition, and no cryptographic feature resolves it, because the quantity in dispute is the population rather than the protocol.

Price discrimination is available and it destroys the good it prices. The standard response to two segments with different willingness to pay is to serve them in separate tiers at separate prices, and the tier machinery of Section 22.3 would support it. Applied here it fails in a specific and instructive way.

Separating the segments into tiers separates their anonymity sets, and set size is **the product**. Two pools of half the size are worth less than one pool to every participant in both. Price discrimination on an ordinary good extracts surplus from each segment at its own willingness to pay; price discrimination on a network good extracts surplus while simultaneously reducing the quantity being sold, and here the reduction falls hardest on the segment whose Δ is highest.

> You cannot segment a network good without degrading it. The protocol can price the two tails separately or it can pool them, and pooling is the only configuration in which either tail gets what it is paying for.

The consequence for the fee, stated precisely. Section 13.3.3 binds the fee at the marginal buyer. With two segments sharing one pool, and price discrimination unavailable for the reason above, the protocol is not choosing between Δ_S and Δ_X and it is not collecting their sum. It is setting one fee against whichever segment it is unwilling to lose, and the achievable revenue is *not* additive across the tails:

$$\text{Rev} \;\le\; \big(\min(\Delta_S, \Delta_X) - (c_p - c_b)\big)\cdot(Q_S + Q_X) \quad \text{under pooling, or} \quad \text{separated tails at reduced } \Delta \text{ each.}$$

Neither branch yields the quantity a naive reading of Section 13.3.5 suggests. f **and** Q **are anticorrelated through the same distribution**: a fee set at the higher tail sheds the volume of the lower one, and a fee set to retain both is bounded by the lower. The document has been treating "is the tail large enough" as a sizing question with one answer. It has two, and they interact.

Which architecture serves which segment. Section 8.3 compares three architectures without pre-selecting one, and this subsection supplies a discriminating criterion the comparison did not previously have:

- **Architecture A** (one native monetary asset) forces both segments into one pool, one governance, and one legal posture. It maximizes set size and maximizes the corrosion. It is the configuration in which an enforcement action against Segment X is an enforcement action against Segment S's counterparties.
- **Architecture B** (neutral reserve plus typed service claims) allows the segments to sit in different service claims over a shared reserve. It is the only one of the three in which the pools can differ while the monetary object does not.
- **Architecture C** (shared settlement, modular domain collateral) separates the collateral but not necessarily the sets, so it inherits the corrosion without the containment.

This does not select an architecture. It states, for the first time in this document, what the selection would have to be made *on*: not cross-subsidy and governance scope alone, but whether the architecture can hold two mutually corrosive high-Δ populations without losing one of them.

The instrument, and why it cannot be the obvious one. The obvious measurement is anonymity-set composition by segment. **That measurement is not available and must not be built**, because a set whose composition is published is not an anonymity set — the instrument would destroy the property it measures, and a protocol that could report who is in its pool has already failed Item 9.

What is observable is *departure*, which is a behavioural series rather than a census:

- **Compliance-side attrition after enforcement events.** Supervised counterparty count and volume in the windows following a designation or enforcement action touching the protocol's pools.
- **Exit-side attrition after disclosure events.** Shielded-pool activity and non-custodial share in the windows following the shipping or mandating of any disclosure surface, viewing-key requirement, or perimeter identification.
- **The cross-correlation of the two series** over rolling windows. A persistently negative correlation is the corrosion, measured without anyone being identified.

A protocol that shows both series declining together is losing the tail from both ends, which

is the configuration in which the fee bound of Section 13.3.3 goes to zero without any bypass channel being responsible.

13.3.7 Two Bypass Channels the Thesis Has Not Been Counting

Both of the following erode the base without tripping the instruments currently pointed at bypass.

Partial bypass at the intensive margin. Bypass has been modelled as a routing decision — the buyer uses the protocol or does not. In practice the buyer uses *less protocol per unit of underlying activity*: batching many operations under one proof, aggregating proofs recursively, or using the protocol purely for attestation while executing the work elsewhere. Each is a legitimate engineering optimization, and each cuts native fees per unit of real economic activity. None of it registers as a Red Line 6 style "usage without fees" signal, because usage does still route through the protocol — just thinner. The correct denominator for fee intensity is the underlying economic activity being attested, not the count of protocol interactions, and the thesis has not been measuring it that way.

The protocol can be forked. This is the more important omission. In an open-source, open-admission, canonically-specified market, the cheapest bypass channel is not a hyperscaler. It is *this protocol, copied, with a lower fee and no burn.* Every property the thesis is proud of — published workload specifications, open verifier implementations, permissionless participation, no vendor chokepoints — lowers the cost of producing that copy. A fork inherits the technology and discards the tax.

What actually prevents the fork from winning is not any of the five conditions. It is liquidity depth in the native asset, anonymity-set size in the privacy corridors, the installed collateral base and the operator relationships around it, the Schelling-point status of the workload registry, and the accumulated receipt history that makes one chain's attestations worth citing. Those are network effects and coordination assets. They are not enforcement rules, and the thesis has been attributing the moat to enforcement rules. Stated honestly: **Section 13.5 explains why a user cannot get the service without the asset on this protocol. It does not explain why the user cannot get it on the copy.** Whatever answer exists is a liquidity and coordination argument, it is empirical, and it belongs to link 7; whether holders then warehouse the asset's risk belongs to link 8.

This is not purely hypothetical, and the one well-documented case ran exactly this experiment against exactly this kind of moat. In 2020, SushiSwap copied Uniswap's open contracts and paid liquidity providers to migrate — a deliberate attack on a liquidity-and-coordination moat by an identical copy with better short-term incentives. The moat held. Liquidity depth, accumulated positions, and the receipt history a trader's history lives in did not transfer with

the code, and the copy won a share rather than the market. That is one observation in one corner of the design space, not proof — but it is evidence that the moat in question is real, and evidence about *which* moat: what defended the incumbent was not its technology, which was copied freely, but the coordination assets stacked on top of it. It is precisely those assets, and not the open specification, that a fork of this protocol would have to defeat.

13.3.8 The Fork Is an Equilibrium, Not a Competitor

The cheapest bypass is not a hyperscaler. It is this protocol, copied, with the fee removed.

Section 13.3.7 concedes that, and concedes what actually stops the copy: liquidity depth, anonymity-set size, the installed collateral base, the registry everyone cites, the receipt history worth citing. Network effects, in other words, and not one of the five lemma conditions. The concession is correct. It was filed under the wrong heading.

Modelled as a bypass channel, the fork looks like the end of the fee. Modelled correctly, it changes what the fee is bounded by.

Two different bounds, and the binding one is not always the same. The sustainable-fee bound of Section 13.3.3 is $f \leq \Delta - (c_p - c_b)$, where c_b is the all-in cost of the best bypass channel. Against a hyperscaler that expression carries the structurally positive cost handicap the section emphasizes: verified compute costs more to produce than unverified compute, and the protocol must recover the difference out of Δ before collecting anything.

Against a fork the arithmetic is different in a way worth stating explicitly. **The fork runs the same technology, so $c_b = c_p$ and the handicap vanishes.** What remains is

$$f \leq \Delta_{\text{coord}},$$

the differential value of transacting on *this* instance rather than on an identical copy — which is not a cryptographic property at all. It is liquidity depth, anonymity-set size, the installed collateral base and the operator relationships around it, whether the workload registry is the one everybody cites, and whether a receipt from this chain is worth citing because of what else is in its history.

The two bounds are separate quantities and the protocol is subject to both:

$$f \leq \min\Big(\underbrace{\Delta_{\text{tech}} - (c_p - c_b)}_{\text{versus the unverified alternative}} , \underbrace{\Delta_{\text{coord}}}_{\text{versus the identical copy}} \Big).$$

This is a real tightening of Section 13.3.3, which stated only the first term, and it is the correct place to put the concession the previous subsection made informally.

Contestable markets and coordination games give opposite predictions, and only one of them is the right model. The objection that opened Section 13.3 — a contestable market competes economic profit to zero — is the right model for the hyperscaler comparison and the wrong one for the fork. Contestability assumes entrants and incumbents are substitutes at the margin for each buyer independently. They are not, when the good's value to each buyer depends on how many other buyers chose the same instance. That is a coordination game, it has multiple equilibria, and equilibria are not competed away by a better price. They are displaced by a focal point moving, which is a different and much slower mechanism.

> A contestable market has one equilibrium and prices race toward it. A coordination game has several, and the incumbent's advantage is not that it is better but that everyone is already there.

Openness cuts both ways, and the thesis has counted only one direction. The previous subsection observed that every property the thesis is proud of — published workload specifications, open verifier implementations, permissionless participation — lowers the cost of producing the copy. True, and incomplete. Those same properties operate on the two bounds in *opposite* directions:

- Openness drives Δ_{tech} against a fork to zero. There is no technical differentiator, by construction and by design commitment.
- Openness *raises* Δ_{coord}. A canonical workload registry that is public, stable, and cited by everyone is more focal than a proprietary one, not less. Schelling points are made of common knowledge, and publishing is how common knowledge is manufactured. The same is true of an open receipt history: its citability is a function of how many independent parties already read it.

So the honest statement of the trade-off, in the form Section 13.3.12 uses for collateral:

> **Openness must be total for the technology and is the mechanism by which the coordination assets accumulate.** You cannot have a defensible technical moat and an open specification. You can have an open specification and a focal point, and the second is the only moat this design is entitled to.

And the moat is buyable, which is the part that keeps it honest. Δ_{coord} is not a law of nature. A challenger can rent it: subsidize liquidity, pay operators to migrate, underwrite the anonymity set, and buy the switching costs down. This is the fork's actual attack and it is a capital question rather than a technical one. The defensible claim is therefore narrow: **the fee is bounded by what it costs a challenger to move the focal point, and that quantity is empirical, contestable by capital, and not guaranteed to exceed the fee.**

The one case, in the right reference class and with its limits stated. Section 13.3.7 cites the 2020 SushiSwap attack on Uniswap: an identical copy of open contracts, paying liquidity providers to migrate, attacking a liquidity-and-coordination moat directly, and winning a share rather than the market.

Two corrections to how that case was used. First, the relevant reference class is not "forks" in general but *coordination games in which an installed focal point is challenged by an identical copy with better short-term incentives* — and in that class the observation is on point rather than cherry-picked, because it is a direct test of Δ_{coord} against paid switching. Second, and stated plainly because the thesis holds itself to base rates elsewhere: **it is one observation, it was selected after the fact, the incumbent survived, and a single surviving incumbent is not a base rate.** It establishes that coordination assets *can* hold against a funded identical copy. It does not establish how often they do, or that they would here.

What this changes in the instruments. Two additions, both small and both closing gaps that the bypass framing left open:

- **Red Line 14's reference panel gains a fork row.** The pre-committed panel of Section 30.6 lists hyperscalers and GPU marketplaces — venues that measure Δ_{tech} only. The lowest-fee protocol-grade fork offering the same canonical workload must be priced into P^{byp} on the same terms, because it is the channel with no cost handicap and therefore the one that binds first. A wedge measured only against unverified alternatives is measured against the weaker competitor.
- **Acceptance breadth becomes a published series.** The number of independent venues, corridors, and counterparties that quote, settle, and cite this instance rather than any other is the observable form of Δ_{coord}, and it belongs beside the denomination series of Section 13.7.3 on the Value Capture Board (Section 26.9). Both measure coordination; one measures whether people reckon in the unit and the other whether they converge on this instance of it.

Where this leaves link 7. Non-bypassability was reduced from Strong to Medium in Table 8.1 because the economics requires a magnitude that is nowhere estimated. Nothing here raises it. The magnitude is now specified as *two* magnitudes rather than one, the second of which the document had been treating as an admission rather than a term, and both remain unestimated. What has changed is that there is now an instrument pointed at each.

13.3.9 What the Five Conditions Establish, and What They Do Not

Collecting the above fixes what the lemma's conditions can deliver.

What they establish. That a contestable market can be taxed and the proceeds routed to a token. Specifically: (i) a competitively priced claim on a cash-flow stream $f \cdot Q$, bounded above by the per-unit wedge Δ applied to volume — $\big(\Delta - (c_p - c_b)\big) \cdot Q$ at the sustainable-fee bound of Section 13.3.3 — and deflating with unit price; and (ii) a balance-sheet lockup of collateral required in the native unit. Both are real. Both are also, precisely, quantities that a discounted-cash-flow valuation captures — a fee stream and a lockup.

What they do not establish. Monetary premium. A monetary premium is by construction value *in excess of* discounted cash flows; it is what gold has and a pipeline does not. A holder who buys a share of the fee stream from indispensable workloads owns something closer to equity in a utility than to a collectible, and a utility equity earns its cost of capital; it does not carry monetary premium. The routing rules of Section 13.2 are design commitments; nothing in this thesis claims that they set a price.

Three corollaries follow:

- **A burn is a buyback.** Economically, retiring supply with revenue is a pro-rata return of value to an existing claim. Buybacks change per-share value; they do not confer moneyness. No quantity of buyback has ever made an equity a store of value, and no burn schedule will do it either.
- **Collateral is a lockup, not a premium.** A required lockup denominated in the asset immobilizes float while operators remain solvent and choose to stay, and it places no bid. It is a balance-sheet fact, and Section 13.3.12 bounds how large a fact it can be.
- **A subsidy is a negative fee.** Pass [Pass 2026] models proof-of-useful-work economies as three regimes: one in which the work is worthless and the token is pure seigniorage, one in which useful work is sold at cost, and one in which token issuance subsidizes the work below its compute cost. The fee-incidence bound places the second regime at $f \leq \Delta - (c_p - c_b)$ and the third at $f < 0$. A network in the third regime is buying volume with issuance; its receipts are real and its fee coverage is negative, and Red Line 5 and Red Line 14(B) are written to see it. Bitcoin's work function, which has no external buyer, sits in the first regime by construction, and that is the content of the concession in Section 34.1.1.

The Test That This Reasoning Does Not Prove Too Much

If durable fee extraction from a commoditized flow made an asset money, the world would be full of monetary assets. CME clears standardized futures; DTCC settles equities; Visa takes a spread on card turnover; MSCI licenses indices that everyone benchmarks to; pipelines and port authorities meter physical throughput. All of them sit on flows more commoditized than canonical MatMul, all of them extract fees durably, several are structurally harder to bypass than any protocol will be, and *not one of their equities is money.* Visa's take rate does not make Visa shares a store of value; it makes

them a good business.
The comparison is the discipline. If the argument for the asset's monetary premium is "recurring fees on indispensable throughput," the argument is an argument for a high-quality utility equity, and it should be labelled as one. Something else has to be doing the monetary work.

13.3.10 Pledgeability Without Due Diligence

Nobody examines a banknote. That is not carelessness; it is the property that makes it money. A thing accepted only after inspection is being traded; a thing accepted without inspection is being spent, and the distance between those two — rather than any fact about scarcity or cash flow — is what the collateral literature means by moneyness.

This thesis has been engineering that property for six Parts and calling it a cost metric. Section 13.3.9 established what the five conditions do not reach; Section 13.3.11 relocates the claim to a holder-side flow. Between them sits the mechanism that arrives first and measures soonest.

The property that the collateral literature actually calls moneyness. Outside the store-of-value framing, the working definition of a money-like asset in the study of collateral and short-term funding is not scarcity, not cash flow, and not convenience yield. It is **information-insensitivity**. An asset is money-like when no counterparty gains anything by investigating the particular unit tendered. Nobody looks, so it passes at face value, and a thing that passes at face value can be pledged without due diligence [Gorton & Pennacchi 1990] [Dang et al. 2020] [Holmström 2015]. Deposits, bills, and repo collateral circulate on that property. It is the property whose failure, rather than any change in fundamentals, is what a funding run consists of.

Two kinds of information, and only one of them is at issue. The claim collapses immediately unless the object is specified, because transparency plainly does *not* produce insensitivity in general—equities are the most transparent securities in existence and among the most information-sensitive. The distinction that resolves this is between:

- **Fundamental information**—what the asset will be worth. Cheap verification does nothing to this, and this thesis claims nothing about it. A triad asset would remain fully information-sensitive on fundamentals, and so does gold.
- **Unit-quality information**—whether *this* unit, *this* receipt, *this* pledged position is what it purports to be. This is authentication and due diligence on the collateral itself rather than a forecast of its value.

The 2008 collapse in securitized funding is the reference case and it was the second kind. The question that stopped the market was not what housing would do; every participant held a view on that and had held one for years. It was whether a *particular* pool was good, nobody could answer it cheaply, and so lenders began producing private information about specific collateral, haircuts widened, and no-questions-asked acceptance ended [Gorton & Metrick 2012]. Moneyness failed along the unit-quality axis while the fundamental axis was already common knowledge.

Information-Insensitivity by Symmetric Knowledge

There are two routes to no-questions-asked acceptance, and they are opposites.

Symmetric ignorance is the historical route. Producing information about the specific collateral is expensive relative to the gain, so nobody produces it, and the asset trades at par. Opacity is the mechanism, not a defect of it.

Symmetric knowledge is the route a verification-native asset takes. Producing that information costs approximately nothing *and the result is public*, so no counterparty can come to hold an informational edge over another. The asset trades at par for the opposite reason.

Both deliver acceptance without due diligence. They differ in how they end. Symmetric ignorance is a knife-edge: it holds only while nobody looks, and the first agent to look acquires an advantage that every other agent must then defend against, which is the run. Symmetric knowledge has no such trigger, because there is no private information *about unit quality* available to produce; private views about fundamentals remain, and are priced as volatility rather than as due diligence. **You cannot start a run by looking when looking is free and public, because looking buys nothing that everyone else could not have had for the same nothing.**

The mechanism is deterrence, not audit, and the distinction is the one a critic will press first. The formulation above deliberately does not claim that everyone has in fact looked. Universal audit is observed nowhere, including in the settings where looking is already close to free: public filings sit unread, open ledgers go unexamined, and the record is routinely both complete and unanalyzed — none of it hidden, and none of it examined ([Green 2026f], where the observation and its limits are attributed). An argument requiring every counterparty to have actually verified would be refuted by ordinary observation on its first day, and a version of this section that rested on one would deserve to be.

Nothing in the mechanism requires it. What removes the run trigger is the absence of a *discoverable edge*, and free public verification removes that whether or not anyone exercises it. The information is non-rival and non-excludable, so the agent who looks acquires nothing the agent who declines to look could not have had at the same cost. Symmetric ignorance fails

because looking is privately profitable to the first mover; symmetric knowledge holds because looking is privately worthless to everyone—not because it is universally performed.

The discipline therefore operates on the *claimant*, before the claim is made, in the way the credible possibility of entry disciplines conduct in a contestable market without entry being observed [Baumol 1982]. A misrepresentation any counterparty could falsify at zero cost is not worth attempting, so it is not attempted, so nobody has to check. **The equilibrium described here is one in which verification is constantly available and almost never used, and those two facts are the same fact.** Two consequences follow that matter for how this thesis is read. The first is that the object to measure is the *cost* of verification and not the observed rate of it: Red Line 1 reads the standing threat, and a decline in measured verification activity is not by itself evidence against the mechanism provided verification cost has not moved. The second is that a system can be perfectly auditable and monetarily inert, because deterrence requires the claimant to believe that being caught would cost them something—a condition this thesis instruments separately, and takes up in Section 13.2 and Red Line 14 rather than assuming.

This thesis has been describing that mechanism as a cost metric. The touchstone image of Red Line 1—expensive to smelt, cheap to assay—is not an analogy about operational efficiency. It is a statement that gold is information-insensitive on unit quality, which is why any counterparty could accept a nugget from a stranger without investigating the stranger. Gold's monetary career is better explained by that property than by a convenience yield it does not have (Section 13.3.11 notes gold's storage-theory convenience yield is near zero). The same substitution runs through the rest of the document: $r(W)$ (Section 22.2), the PIDL receipt schema, the reproducibility contract, and the Layer 0 grade ladder are not merely cheaper ways to run an audit. They are machinery for removing the informational advantage that a counterparty could otherwise hold over the holder of a pledged position.

The clearest evidence that the thesis already believes this is that it already prices it. The collateral haircut schedule of Section 17.4 discounts a position by how legible its physical substrate is; risk haircuts (Section 25.6.5) discount it by sovereign optionality; bridge classes carry haircuts by trust assumption. Every one of those is a price on residual unit-quality uncertainty. **The document derived an information-sensitivity schedule from engineering intuition without the monetary theory that explains what it was pricing.**

What is made insensitive is the position, not the token. A necessary precision, because the obvious version of this claim is trivial. That a unit of a public ledger is authentic is not a differentiator: it is true of every cryptocurrency and no counterparty has ever had to investigate whether a bitcoin is a bitcoin. That is table stakes and it has never made a token money.

The differentiator is the *collateralized position*. When the asset is pledged against an obligation to deliver a service—prover bonds, corridor collateral, SLA escrow—what a lender

must assess is not the unit but the bundle: whether the capacity behind the obligation exists, whether the operator's slashing exposure is what it claims, whether the physical substrate can deliver under the scenario the haircut assumes. **That** is where due diligence lives, and that is what Facility Capacity Receipts, energy receipts, Delivered Verified Capacity (Section 17.9) and the grade ladder make cheap. It is the first argument in this document for why Layer 0 is a *monetary* requirement rather than an engineering preference, which is the role Chapter 4 assigns it.

Necessary, not sufficient—and the test that says so. The discipline of Section 13.3.9 applies to this section with equal force, and the counterexamples are the thesis's own. Graded commodities are cheaply assayed and warehouse receipts never became money. Certified diamonds are authenticated to a published standard and are not money. Bearer bonds are authenticable at sight and are not money. Cheap unit verification is therefore not sufficient, and any argument that stops here proves far too much.

What separates the cases is a conjunction. Insensitivity on unit quality becomes monetary only alongside **fungibility** (certification tells you this stone's grade; it does not make stones interchangeable), **durability and portability** (grain is gradeable, perishable, and bulky), and **the absence of a counterparty** (a warehouse receipt reintroduces the custodian whose solvency must then be assessed, and a bearer bond is somebody's liability). Gold is the one historical object at that intersection: assayable, fungible, durable, portable, and nobody's promise. The claim available here is narrow and it is the whole of what this section asserts:

> A verification-native digital bearer asset is the second class of object to occupy that intersection—and the first whose assay cost does not scale with the size of the claim being assayed.

That last clause is the only genuinely new term. Assaying a larger quantity of metal costs more; destructive sampling scales with the lot. Verifying a succinct proof costs the same whether the claim underneath it is worth a dollar or a billion. The cost of establishing unit quality has been invariant to notional in no prior monetary object, and if information-insensitivity is what makes collateral money-like, an assay whose cost does not grow with what is at stake is a change in kind rather than in degree.

What this predicts, and what it does not. A reader of Dang, Gorton and Holmström [Dang et al. 2020] will object that their insensitivity is a property of the *payoff structure*. A security is insensitive when the expected gain from producing private information about its payoff is below the cost of producing it. That is why over-collateralized debt—flat over the relevant states—is the natural money-like security, and why driving information cost to zero on a security with a state-dependent payoff makes it *more* sensitive, not less. The objection is correct,

and it fixes the scope of this section. The argument made here is the adverse-selection argument on unit quality, not a claim about payoffs: free public verification removes the class of private information that concerns whether the pledged position is what it purports to be, and leaves the class that concerns what it will be worth untouched. Bitcoin is the control case. Its unit quality has been perfectly and freely verifiable since 2009, and its collateral haircuts nonetheless sit at 30–50% and vary widely across lenders, because those haircuts are pricing realized volatility and liquidation depth, which no verifier can compress. Cheap verification therefore cannot be expected to lower haircut *levels* on a volatile bearer asset, and the thesis does not predict that it will.

What it predicts is narrower and is the content of Red Line 17. Controlling for realized volatility and market depth, a position whose backing is publicly verifiable at VerifyPrice should show (i) lower *dispersion* of haircuts across lenders than an opaque comparable, because the due-diligence component that lenders price idiosyncratically has been removed, and (ii) faster convergence of haircuts after a shock, because there is no private information for the cautious lender to be waiting on. Pledgeability without due diligence attaches to the *position*—over-collateralized, DVC-backed, with a debt-like payoff over the states the haircut assumes—and not to the bare base asset; a lender who wants a flat payoff builds one out of the base asset and a haircut, and what Layer 0 sells that lender is the ability to size the haircut from public facts. Red Line 17 tests exactly that, and is specified so that Bitcoin's own record—native units against wrapped and custodial representations of the same units—is a case it can pass or fail rather than one it fails by construction.

A second invariance, and the objection it answers. The strongest available objection to any cheap-verification argument is that verification is not the only thing getting cheaper. If the cost of manufacturing plausible false claims falls at least as fast as the cost of checking them—synthetic corroboration, fabricated provenance, volume sufficient to exhaust the checker's attention—then the informational balance is unchanged and nothing monetary follows. The objection is correct in general and it is fatal to arguments about verification conducted in natural language, where discrimination between a true and a false claim degrades as the adversary invests.

Cryptographic verification does not have that gradient, and this is the second respect in which it is unlike every prior assay. Under stated soundness assumptions — the hardness and parameter assumptions, and a circuit and verifier that constrain what they claim — a false proof is not *expensive* to produce; it is infeasible, and the verifier's cost of rejecting one is identical to its cost of accepting a valid one. Attempts do not accumulate into persuasion: a thousand forged proofs are a thousand rejections, and no quantity of surrounding corroboration moves the outcome, because nothing outside the proof is consulted. **Verification cost is invariant to the size of the claim and invariant to the effort of the adversary, and the second**

invariance is the one that survives contact with an opponent who is also getting better tools.

The bound on that claim is stated with equal force, and it is the bound already set out at the head of Chapter 10: a proof is a bounded attestation, not truth. It establishes that a computation followed a circuit or that a receipt descends from a signed chain. It establishes nothing about whether the workload was worth running, whether the operator's off-ledger position is what they say, or whether the sentence a human attached to the artifact is honest. The class of claims over which this thesis's verification is adversarially invariant is exactly the class that can be stated as a computation, and that class is considerably narrower than the class of claims on which trust is currently failing. Where the two do not overlap, this thesis has nothing to offer and should not be read as offering it.

Issuance discipline reappears here, and not as scarcity. Condition 4 of the Value Capture Lemma required issuance to be capped by schedule or capacity rather than by governance fiat, and Section 13.3.11 observes that issuance discipline is the one lemma condition touching the monetary channel without explaining what distinguishes it. This section supplies the explanation. An asset whose supply is a discretionary decision is information-sensitive on supply: a counterparty must form and maintain a private view of what the issuing body will do, which is due diligence by another name, and an asset requiring due diligence cannot be accepted no questions asked. **Credible non-discretion is not only a scarcity property. It is the condition under which the supply schedule stops being something anyone has to have an opinion about.**

Who pays for verification, and why the free-rider problem does not bind here. Grossman and Stiglitz [Grossman & Stiglitz 1980] established that a market cannot be both costlessly informative and informative at all: if prices already reflect everything knowable, no agent is compensated for the costly work of producing information, so the work is not done and the prices stop reflecting it. Informativeness has to be paid for by somebody, and an equilibrium in which nobody is paid is not an equilibrium. The same structure threatens the argument of this section in an obvious form. If verification is free, nobody earns a return on verifying; if nobody earns a return on verifying, what sustains the apparatus that makes verification possible, and why does the symmetric-knowledge regime not decay back into the ignorance it replaced?

The objection is a real one, and it is the reason this section does not rest on transparency in the general case—where it would bind, and where the historical record of unread disclosure suggests it does. What answers it is a property specific to proof systems and absent from the Grossman–Stiglitz setting: **the cost of producing the information and the benefit of consuming it fall on different parties, and the party bearing the cost bears it voluntarily.** In their model the same act, research, is simultaneously the expenditure and the source of the private return, which is what forces the informational rent that makes prices imperfectly informative.

In a proof system the expenditure sits with the *prover*, who incurs it in order to be believed, and the result is consumed by every counterparty at negligible marginal cost. Verification is therefore not an underprovided public good waiting for a subsidy or a rent. It is a private good purchased by the party who wants their claim accepted, and its public checkability is a by-product that the purchaser cannot withhold without destroying the thing they were buying.

The same asymmetry, read from the financing side, is why the fee argument had to be conceded. Section 22.2 states this separation as $r(W) = v(W)/p(W)$ and treats it as an engineering property. It is also an industrial-organization property, and reading it that way closes a loop this Part opened. Because the prover is buying credibility in a market where any competent competitor can sell the same credibility, the price of proving is competed toward the cost of proving. That is precisely the finding of Section 13.3.9: the fee stream is ordinary service revenue, competitively priced, and reproducible by a discounted-cash-flow valuation. **The property that makes verification sustainably supplied is the same property that stops it from being a source of excess return, and the thesis should not have expected to keep both.** The two results are one result seen from opposite sides, and their agreement is a constraint the argument has to live inside rather than a coincidence worth celebrating: whatever monetary premium exists is not in the proving business, which is why Section 13.3.11 and Section 13.7.3 look for it somewhere a valuation cannot follow.

How symmetric knowledge fails, stated before a critic states it. The mechanism is not indestructible, and its failure modes are already red lines, which is the useful part. Symmetric knowledge reverts to ignorance whenever verification stops being free and universal: when verification affordability breaks (Red Line 1), when verification concentrates such that most participants are trusting a verifier rather than verifying (Red Line 3), or when the underlying receipt data becomes unavailable or proprietary (Red Line 4). Each of those has been treated in this document as a degradation of *service*. Each is also, on the argument above, a **monetary run trigger**, and should be read as one.

A fourth failure mode is not covered by any red line and is named here rather than deferred: **scope drift between what is verified and what is exposed**. A position can be perfectly legible on its proof and entirely opaque on the operator's off-chain leverage, counterparty concentration, or rehypothecation. Verification that certifies the computation while the exposure sits elsewhere reproduces exactly the 2008 structure with better cryptography. The gap is measured as the **off-ledger exposure disclosure** series on the Value Capture Board (Section 26.9, Appendix B): the share of pledged-position notional whose operator publishes an attested disclosure of off-ledger leverage, counterparty concentration, and rehypothecation, bound to the position's receipts. It carries no trigger of its own in this version; it feeds the RL17 lender panel

and the RL10 opacity reading, and it exists so that the gap is a published number rather than a caveat.

What this section does not claim. Four disclaimers, because the temptation to over-read this is considerable. It does not claim price stability: an information-insensitive asset can be violently volatile, and gold and Treasuries both are. It does not claim insensitivity on fundamentals, only on unit quality and pledged-position quality. It does not claim that acceptance follows automatically from measurability; whether any lender actually extends no-questions-asked terms is an empirical outcome, not a design consequence. And it does not raise any rating in Table 8.1: this section supplies a mechanism and an instrument for link 9, not evidence that link 9 holds. Ratings that rise because an argument was added are self-congratulation.

The instrument: haircut level, and more sharply, haircut dispersion. The mechanism is measurable well before the asset is anything, and one of the two available readings is much cleaner than the other. Haircut *level* confounds what this section is about with ordinary price volatility: a lender demands over-collateralization for both reasons and the level does not separate them. Haircut *dispersion across lenders* separates them better, but not by itself: lenders differ in funding costs, balance-sheet limits, liquidation capacity, and terms, and lenders running one shared model can agree while sharing its error. The comparison that isolates the sensitivity at issue is relative — dispersion on the native asset against dispersion on an opaque comparable of similar volatility, quoted by the same lenders at the same moment — because whatever is particular to each lender enters both quotes and differences out. What remains is evidence about unit-quality information under that comparison, not an identity with it, and agreement produced by a shared underwriting model is not counted (Red Line 17).

Publish, per quarter, against a pre-committed panel of lending venues on the same custody and reproducibility terms as the Red Line 14 reference panel:

- **Haircut level** on the base asset and on pledged positions, by collateral class and Layer 0 grade.
- **Haircut dispersion**—the cross-lender interquartile range at a common timestamp—as the primary series.
- **Due-diligence cost**: the observable cost of onboarding the asset as collateral at a new venue, in time and fees, which should fall toward zero if the mechanism is working.
- **Dispersion conditioned on stress**, since the prediction with the most content is about the shape of failure rather than its level.

The prediction is specific enough to lose. An information-insensitive asset shows haircuts that are *tight* at whatever level volatility sets, and—this is the discriminating claim—dispersion that does not widen when the market does, because there is no private information for a stressed lender to price. The historical signature of a funding run is dispersion widening

before level, as the better-informed lenders withdraw first. If that signature appears here with full receipt availability, symmetric knowledge did not hold, and the mechanism this section proposes is not present in the asset.

The falsifier. This mechanism is made falsifiable as **Red Line 17** (Section 30.6): cross-lender haircut dispersion that fails to decline as verification coverage rises, or that widens under stress while receipts remain available, is this section's claim failing in public.

Two mechanisms, and they are complements. This section and Section 13.3.11 are not rival accounts, and the difference between them is worth stating in one line because they answer different questions. Information- insensitivity is why a counterparty would *accept* the asset without investigating it. The regime-contingent holder-side flow is why a holder would *want* it when other arrangements fail. Acceptance without demand is a payment rail; demand without acceptance is a collectible. A monetary asset requires both, and Section 13.7.3 adds the third question — whether anyone will *reckon* in it. The thesis argues all three, sizes none, and publishes an instrument for each.

13.3.11 Where a Monetary Premium Would Actually Come From

The constructive half. If fees, burns, and collateral yield a well-valued cash-flow claim rather than money, the monetary claim has to be relocated rather than abandoned — and the place it relocates to is one the thesis has already spent six Parts describing without noticing that it was the monetary argument.

Δ is not a constant; it is a function of regime pressure. Write $\Delta(R)$, with R the regime-pressure index of Stage A (Section 13.7.1). Section 13.3.5 showed why Δ is thin in benign states: a solvent, suable indemnitor substitutes for a proof. Invert that sentence and it names exactly when Δ is large. The substitute fails when the counterparty:

- **cannot be sued** — cross-jurisdictional, pseudonymous, or sanctioned, so no forum will hear the claim;
- **will not be solvent when it matters** — the failure is correlated, and the indemnity is worth least in exactly the state that triggers it;
- **is the adversary** — a state, a platform, or a regulator, where the entity you would sue is the entity imposing the loss;
- **inflicts a non-compensable loss** — censorship, seizure, deplatforming, or disclosure, where damages do not restore the position because the position was never about money.

That list is the threat model of Chapter 5, restated as a demand curve. The thesis is therefore internally consistent — but the consequence is unflattering to fee accrual as a monetary argument and places the monetary burden on links 8 and 9:

> Δ is small in benign states and large in the tail. The asset's usefulness is state-contingent, and its distribution is skewed toward exactly the states everyone else's arrangements stop working in.

State-contingency does two different things, and only one of them is monetary. The tempting move here is to say that an asset whose usefulness rises in states where the marginal utility of wealth is high must be worth more than the discounted expectation of its cash flows, because those cash flows arrive weighted by a high state price. That move fails the test this section's own demolition just set. Section 13.3.9 defined a monetary premium as value *in excess of* discounted cash flows, and a discounted-cash-flow valuation carried out with the correct stochastic discount factor already weights every cash flow by its state price. State-price weighting is what a risk-adjusted discount rate *summarizes* — and a single discount rate summarizes it correctly only when payoffs are not heavily state-contingent, which this asset's are. A payoff this state-contingent needs the stochastic discount factor itself, not the single scalar that compresses it. An argument that stops there has relocated the monetary claim into the denominator of a DCF, which is not outside the DCF at all. The argument therefore has to be split into two mechanisms, one conceded and one kept.

Mechanism one: a negative beta on the fee stream. Real, valuable, and not monetary. Because $\Delta(R)$ rises when the substitutes fail, the fee stream is countercyclical with respect to regime pressure: the protocol can charge most for its differential properties in the states where wealth is scarcest. A rational investor pays more per unit of *expected* cash flow for that stream than for a procyclical one, and should. But this is exactly a discount-rate effect. It lowers the cost of capital applied to the same fee stream, it is fully representable inside a discounted-cash-flow valuation, and it therefore produces no moneyness. A countercyclical utility is a better utility, not a monetary asset. Conceded, and set aside: nothing in the monetary claim below rests on it.

Mechanism two: a holder-side service flow. This is where a monetary premium would have to live. The bearer of the asset obtains something that is not a cash flow at all, and so has no cash flows to discount: the ability to transact, prove, hold, and exit at the moment when the substitutes for proofs, courts, and custodians have stopped working. That service accrues to whoever holds the unit, by virtue of holding it, and not by virtue of any distribution from the protocol. Nothing is paid out, so a valuation of distributions omits it however its discount rate is set; a valuation that prices services from possession, as transactional-demand models of tokens do, need not. This is the same category of thing as the value of holding a bearer instrument outside the banking system, and it is why an asset with no cash flows whatsoever can carry a price. It establishes a benefit of ownership. Whether that benefit becomes monetary

use rather than inventory demand depends on others accepting the unit, which is the separate question of Section 13.3.10.

The mechanism is standard, and the analogy is to the mechanism only. That a non-cash-flow service from *holding* a thing can carry a price no cash-flow model reproduces is the storage-theory account of **convenience yield** in consumption commodities ([Kaldor 1939]; [Working 1949]): the holder of physical inventory obtains an option against stockout that the holder of a claim on future delivery does not. The analogy drawn here is to that mechanism — a service flow accruing to the bearer rather than to a claimant — and not to the setting. A triad asset is not a storage commodity, it faces no stockout risk in the storage-theory sense, and no magnitude in this thesis derives from the commodity literature.

The distinction the mechanism turns on is easiest to see with a fire extinguisher and a contract for the delivery of one. Both have a price. Both can be bought today, valued, resold, and hedged. Only one of them works during the fire. The premium you pay for owning the extinguisher rather than a delivery claim on one is the amount by which the physical commands more than the discounted contract. No discounted-contract valuation reproduces that number. It is the quantity this thesis means when it says the holder-side flow has no cash flows to discount. In the moment of use, the contract and the extinguisher are not substitutes at any price, and every argument in this section about regime states is an argument about when the market is on fire. An extinguisher is not money, though. The analogy establishes a benefit of ownership; the step from there to monetary use is independent acceptance and a persistent demand to hold, and this section does not supply it.

Gold, which this thesis keeps comparing itself to, is a separate case and should not be filed under storage theory. Its convenience yield in the storage-theory sense is near zero in normal regimes, which is why its lease rates are ordinarily negligible and its forward curve essentially full cost of carry. The early-2025 episode, when tariff-driven physical flows into New York pushed one-month lease rates to multi-year highs and spot into backwardation, is the exception that shows what a storage-theory yield on gold looks like when it briefly exists. What is true of gold is the part that matters here. It has no cash flows, negative carry, and industrial demand far too small to explain its price; every constructible DCF for gold returns a number well below the market. On the standard account that residual is monetary and portfolio demand — what holders pay for a bearer holding that is useful when other arrangements are not — which is the holder-side category described above, arrived at from a different direction.

Regime-Contingent Convenience Yield

The monetary premium available to a triad asset *is* the **holder-side service flow** generated by the state-contingency of $\Delta(R)$ — the bearer's ability to transact, prove, hold,

and exit when the substitutes for proofs, courts, and custodians have stopped working — *conditioned on* credible non-discretion in issuance and on bearer holdability. Those two are not added terms but enabling conditions: the flow is only a monetary service while issuance discipline holds and the unit remains bearer-holdable, and it collapses if either fails. The service accrues by virtue of holding rather than as a distribution, so there is no stream to discount. That is what places it outside a discounted-cash-flow valuation rather than merely at a different discount rate inside one.

It is *not* the countercyclicality of the fee stream. That is a negative beta on cash flows: real, valuable, and captured exactly by a DCF using the correct stochastic discount factor, which is why it confers no moneyness.

It is also not generated by burns, by collateral, or by fee share. Those produce a cash-flow claim and a collateral lockup (Section 13.3.9).

Which conditions and which links this actually implicates. The consequence of the split is immediate. If the monetary claim rests on the holder-side flow, then it rests on the properties that determine whether a holder can actually use the thing under pressure — and none of those are fee-routing properties. Of the five lemma conditions, only **issuance discipline** touches the channel at all, and it does so because credible non-discretion is part of the payoff rather than part of the accrual. The other four are conditions for value *accrual*, and accrual is a DCF quantity. Link 8 then requires a persistent self-custodied constituency capable of bearing loss. The premium itself hangs off link 9: liquidity, neutrality, verifiability, legal holdability, and agency preservation (Item 9). Those are the conditions under which a bearer can transact, prove, hold, and exit when it counts. No burn schedule reaches the premium: a burn operates on the claim, and the premium is not on the claim.

Now discount it honestly. The hedge has a defect, and it is the same defect the thesis has already documented without drawing the monetary conclusion. The payoff is negatively correlated with the states you fear — that is the point — but the *delivery* is positively correlated with them. The states in which $\Delta(R)$ spikes are energy interdiction, hardware denial, network filtering, and jurisdictional pressure (Section 5.5, Section 17.10.1). Those are precisely the states in which the protocol's ability to supply privacy, proofs, and compute is impaired. Section 34.1.7 half-concedes this and Red Line 13 makes the mechanism falsifiable, but both frame it as a threat to the hinge. It is also a *discount on the hedging premium*, which is the monetary consequence, and the honest statement of it is: you are being asked to pay a premium for insurance whose underwriter is exposed to the insured event.

Two further deductions from the same discount:

- **The slots are occupied.** Gold already holds "works when the network does not." Bitcoin

already holds "works when the state does not." Neither requires a grid at the moment of use — gold requires none at all, and a Bitcoin key survives an outage even where settlement waits. A triad asset needs power, reachable networks, and obtainable reference hardware in order to deliver anything at all, which makes it a hedge against a narrower and more specific band of failure: adversarial conditions severe enough to need proofs and privacy, but not severe enough to take the grid.

- **The band may still be the important one.** That band — administrative repression, surveillance, synthetic media, platform and compute enclosure, financial repression without collapse — is the one Part I argues is most likely. Sovereign optionality (Section 17.8) exists to widen it, and disruption-adjusted VERIFYPRICE (Section 17.10.2) exists to price whether the widening is real.

And do not overclaim it. The state-contingency argument establishes a *mechanism*. Non-cash-flow service flows accruing to a bearer are ordinary in asset pricing, and there is nothing exotic in claiming one here. The **holder-side flow is nonetheless unsized**: nothing in this document estimates how large a regime-contingent convenience yield a triad asset could command, and no honest estimate is available before there is a market in which to observe one. This is an empirical question the thesis has not answered. What it can do is refuse to smuggle the answer in. That refusal is why the chain-strength rating for the store-of-value premium is Medium (conditional), why link 8 asks who bears the asset's loss, and why the relocated claim is attributed to link 9 rather than to fee accrual. The measurement instrument exists so that the question can be settled rather than asserted: Section 26.8.6 publishes the lending-rate, basis, and wrapper-basis series against the regime-pressure index, and a series flat across regime states is the null result, reported as such.

13.3.12 Collateral: Unit-Elasticity, Float, and Wrong-Way Risk

Condition 3 requires operators to post the native asset as collateral, and the reference designs specify the requirement as a share of *capacity value*: "10–20% of capacity value" in Section 8.4 and a "10–15% collateral requirement" in Design A. That specification has a consequence.

Collateral Unit-Elasticity

Let K be the USD value of deployed capacity, χ the required collateral ratio, and P the asset price. The number of units that must be locked is

$$N = \frac{\chi K}{P}.$$

The requirement fixes a *value*, not a quantity. Collateral demand is therefore unit-elastic in price: a doubling of P halves the units required to satisfy the same obligation.

(χ is used here for the collateral ratio because κ already denotes net mechanical gain in Appendix H and facility capacity share in Section 17.8.)

It is a level effect on free float, not a slope on demand. The mechanism immobilizes a fixed *dollar* quantity of float — a level effect on the free float, not a buy order of any kind, since a lockup places no bid — and the valuation it supports grows only as deployed capacity K grows, not as price grows. Describing it as "structural demand that scales with the network" is true only in the sense that it scales with physical build-out, and physical build-out is slow, capital-rationed, and observable.

And K deflates too. Capacity value is the replacement cost of deployed infrastructure. It therefore inherits the same deflation exposure as the fee base in Section 13.3.2: the same generation of throughput is worth fewer dollars of replacement capital each year. A collateral requirement pinned to K shrinks in dollar terms unless deployed physical capacity grows faster than its unit cost falls. This is the same empirical bet, entering through a second door.

Reflexivity runs in both directions. For small moves the mechanism is genuine negative feedback, and it is a real design achievement: price rises, fewer units are needed to cover the same χK, surplus collateral releases into float and meets the buying; price falls, operators must post more units, and they buy. An automatic stabilizer.

For large moves it inverts. Collateral denominated in the asset it insures loses value exactly when slashing risk and operator distress rise — textbook wrong-way risk, and in its dynamic form the familiar margin-spiral or leverage-cycle mechanism. The operator's instruction in a drawdown is "post more units, at a moment when your existing holdings are worth less and your operating business is stressed." A common answer is not to post; it is to *exit capacity*. Exiting reduces K, which reduces required collateral, which releases locked units into a falling market. The stabilizer becomes an accelerant, and it does so through the same equation that made it a stabilizer:

$$P \downarrow \Rightarrow\ N \uparrow \Rightarrow\ \text{operator exit}\ \Rightarrow\ K \downarrow \Rightarrow\ N \downarrow \Rightarrow\ \text{float} \uparrow \Rightarrow\ P \downarrow .$$

The value-leakage mitigations of Section 13.2 do not contemplate this loop. Red Line 15 (Section 30.6) does: it reads on exactly this sequence as a conjunctive condition — a material drawdown, collateral coverage falling below realized slashing exposure, operator exit or forced liquidation, stress-adjusted DVC declining, released collateral raising liquid float, and core SLOs failing to recover — and Section 30.5 requires the pure-native, mixed, and native-plus-insurance designs to be stress-tested against it before governance chooses among them. The

metric the loop needs, collateral coverage against realized slashing exposure with capacity-withdrawal rates conditioned on drawdown, is the one Red Line 15 publishes.

The design tension, stated plainly. Collateral must be the native asset in order to create native demand, which is the whole of Condition 3's monetary contribution. Collateral should be uncorrelated with the risk it insures in order to be credible, which is the whole of prudential collateral design. **You cannot have both.** Any move toward external collateral — stablecoins, bonded fiat, over-collateralization in an uncorrelated asset — strengthens safety and weakens Condition 3's monetary contribution by exactly as much. This is not a problem to be solved; it is a trade-off to be chosen deliberately and disclosed.

The magnitude bound. χK is bounded by an accounting identity, and the arithmetic is worth doing even roughly. The following is an order-of-magnitude illustration, not an estimate: the inputs are chosen to be generous rather than defensible, and no claim rests on their precision.

Deployed capacity value K	**Locked at** $\chi = 0.15$	**Comparable**
\$1T (deliberately generous)	\$150B	A single large-cap equity
\$100B (more defensible)	\$15B	A mid-cap equity

Table 13.3: Order-of-magnitude ceiling on collateral lockup. Illustrative arithmetic on assumed capacity values, not a forecast; the point is the exponent, not the digits.

Those are large-cap-equity numbers. They are not monetary-aggregate numbers, and no choice of χ inside a plausible range changes the exponent. Collateral therefore cannot be the source of moneyness. At best it supports holding demand while operators remain solvent and choose to participate, and it releases in the drawdown that tests it.

The one real contribution, stated in the document's own notation. Collateral does do something, and it is worth being precise about what. It creates a holding constituency whose supply is highly price-inelastic — operators cannot sell locked units at any price without exiting the business — which shrinks effective free float and steepens the price response to a given flow. In the Stage C decomposition (Section 13.7.1), that is an operation on the liquidity and market-impact state Λ_t, not on the native monetary impulse θ_t.

> Collateral is a slope effect on price impact, not a level effect on price. By immobilizing float (a level effect on *free float*), it changes how much price moves per unit of flow. It does not establish where price belongs, because it places no bid.

And a slope effect cuts both ways: a steeper response to buying is a steeper response to

selling. Calling this "structural demand" overstates it in one direction while concealing the symmetry in the other.

13.4 From Utility Demand to Monetary Premium

Many indispensable services—electricity, bandwidth, compute instances, legal services, cloud storage—have recurring demand without being stores of value. The monetary argument requires explaining why *this* utility becomes **monetized savings demand**, not just operating expense.

The bridge has eight conditions. The middle five are the Value Capture Lemma arriving early and informally; they are named here in one line each and stated formally at Lemma 13.1, not argued twice:

1. **Utility demand is not enough.** Recurring demand for a service does not automatically create a scarce, holdable asset. Bandwidth is indispensable, but claims on bandwidth do not become money.
2. **The asset must be required for fees or settlement** — required fee medium, Condition 1.
3. **Fees must produce burns, retirement, or staking yield** — supply reduction, Condition 2.
4. **Operators must post the asset as collateral** — Condition 3, whose worth Section 13.3.12 bounds as a slope effect on price impact, not a valuation floor and not a source of premium.
5. **Issuance must follow a schedule fixed ex ante** — issuance discipline, Condition 4.
6. **Users must not be able to bypass the asset at equal service quality** — non-bypassability, Condition 5.
7. **Telemetry must prove the loop is working.** Fee coverage, burn rates, collateral lockups, and workload demand must be publicly verifiable—not asserted.
8. **If any value-capture condition fails, the asset does not even capture its own service value.** It becomes infrastructure exposure whose economics accrue to operators or customers rather than to holders. Meeting those five conditions makes service value accrue to the asset; it does not make the asset money, establish stress-deliverable service, or create a loss-bearing holder constituency.

This section is the "utility-token trap" defense. The thesis does not claim that useful things automatically become money; it claims that *under specific, testable value-capture conditions*, a useful asset accrues the value of its usefulness to its holders — a cash-flow claim and a collateral lockup. Whether anything further, a monetary premium in excess of those, is available is argued separately and on different grounds in Section 13.3.11. The Value Capture Lemma formalizes the accrual conditions; the telemetry regime (Chapter 25) makes them falsifiable.

13.5 Why Users Cannot Simply Bypass the Asset

The strongest economic objection to the thesis is simple: *"Why can't AWS, a ZK prover marketplace, or a privacy wallet sell the same service for fiat or stablecoins and bypass the base asset entirely?"*

The monetary object at stake in this objection is the **base asset** — the fee, settlement, and collateral unit — not Work Credits. Work Credits are bypassable service claims by design, and their bypass threatens only service-claim economics, not the monetary thesis. What the thesis needs is that the *base asset* cannot be bypassed, and the objection is valid unless the protocol enforces native value capture. The required mitigations:

1. **Core fees are denominated in the base asset.** Provers, routers, and settlement corridors accept only the native asset for protocol-level fees. Fiat or stablecoin payment requires acquiring the asset first — for the seconds it takes to settle, which is transactional demand, not a savings motive.
2. **Collateral must be posted in the base token.** Operators cannot participate without holding significant quantities of the asset. Calling this "structural demand" overstates it in one direction while concealing the symmetry in the other (Section 13.3.12): it immobilizes a fixed *dollar* quantity of float, which steepens the price response to flow in both directions rather than establishing where price belongs.
3. **Fee retirement is a pro-rata buyback tied to usage.** A published share of fees is retired. This returns fee value to holders (Value-Capture Condition 2) and reduces *net* issuance only when retired fees exceed scheduled issuance; it is a cash-flow mechanism and carries no moneyness inference (Section 13.3.9).
4. **Priority is sold in the asset, never held.** Gold-tier SLAs and priority access during congestion are bought through service contracts paid in the asset, which keeps them a fee channel; holding the asset confers none of them (Section 13.1). Governance participation requires staking.
5. **Settlement paths are natively denominated.** Privacy corridors and settlement rails denominate in the base asset; off-ramps exist but are not the default. Whether anyone *keeps* the unit after settling is the denomination question of Red Line 18, not a property of the rail.

Honest admission: If users can in practice bypass the asset—for example, if most operators accept stablecoins and immediately off-ramp, or if hyperscaler-hosted provers dominate the market and set prices in fiat—then the SoV thesis fails. The system may remain useful infrastructure, but the monetary claim collapses. The telemetry must track this: native-asset fee share, bypass channel volume, and operator off-ramp rates should be visible on the Economic Coverage Board (Chapter 26).

13.5.1 The Incumbent Bypass Channel Is Not Hypothetical

The five mitigations above answer a proposal. They do not answer a precedent, and a precedent exists: **stablecoins already settle a large fraction of the value the thesis is targeting**, at scale, at near-zero cost, with no native monetary asset anywhere in the loop. Tron and Ethereum carry stablecoin transfer volumes at national-payment-system scale; in the emerging-market currency episodes Chainalysis documents [Chainalysis 2024] — Argentina, Turkey, Nigeria — dollar-stablecoin wallets, not Bitcoin and not Zcash, were the observed instrument of savings flight and censorship-adjacent settlement. That is a decade of revealed preference in exactly the states of the world this thesis calls State 3 (Section 29.2), and it is not engaged anywhere in the design.

The stablecoin channel is a live test of Condition 5 that has already been run, and the honest reading is unfavorable in three of four respects:

- **Where stablecoins beat the triad: benign and mid-band repression.** A stablecoin settles in seconds, costs cents, has no learning curve, and offers *practical* censorship resistance to the median user, because freezing a Tron address is harder than freezing a bank account even though it is easier than censoring a shielded pool. For payments-motivated flight, the bearer-asset properties this thesis prices (privacy by default, portable proofs, verified compute access) are mostly not on the buyer's list.
- **Where the triad can still differ: the upper band.** The stablecoin's weakness is structural and appears exactly where the thesis's services begin: it is a claim on a custodian's reserve and an issuer's goodwill, it freezes under sanctions enforcement, it carries no proof or compute service, and its censorship resistance is borrowed from its host chain rather than owned. As repression crosses from financial into administrative and epistemic forms — compelled disclosure, platform deplatforming, attestation demands — the stablecoin inherits the substitutability of its issuer, not of cryptography. That is the band in which a triad asset's Δ over a stablecoin is nonzero, and Section 13.3.11 already prices the premium as regime-contingent.
- **The arithmetic.** None of that changes the base fact: for the median settlement need in benign and mid-band states, the incumbent bypass channel is cheaper, faster, and already deployed. The thesis's five mitigations make bypass harder *on this protocol*; they do nothing about bypass *around* it. Red Line 6 is therefore best read as already having an empirical prologue: every stablecoin-settled triad-adjacent payment is a data point that the capturable wedge must be measured against, not assumed above.

The design implication is not despair but precision: the asset's fee-bearing services must be ones a stablecoin cannot express — private settlement where the stablecoin is a custodial IOU, proofs where the stablecoin is silent, verified compute where the stablecoin has nothing to verify. Where the triad service is interchangeable with a stablecoin payment, Condition 5 should be presumed failed, and the Value Capture Board's stablecoin-denominated fee share

series (Section 26.9) is the instrument that watches it fail or hold.

Co-option as Bypass

The bypass threat is not always a competing product. It can also be institutional co-option: ETFs, treasury companies, margin loans, and regulated custody that deliver exposure to the asset's price without requiring users to interact with the protocol's privacy, proof, or settlement rails. If the majority of demand is satisfied by custodial wrappers that bypass the fee-burn-collateral loop, the asset's *price* may rise while its *monetary thesis* weakens. The telemetry must distinguish between custodial exposure (which does not exercise value capture) and protocol-native usage (which does).

13.6 Wrapper Dominance Risk

Co-option as bypass is dangerous enough to warrant a named, top-tier risk rather than a sidebar. A protocol asset can succeed as a financial product while failing as money. If ETFs, custodians, treasury companies, broker-dealers, lending desks, and stablecoin wrappers satisfy most demand for exposure, the asset's price may rise even as native settlement, privacy usage, fee burns, collateral lockups, and non-custodial flows stagnate. This is not a contradiction. It is the signature of **wrapper dominance**.

Wrapper dominance is dangerous because it turns a monetary network into a reference price. It creates holders without users, exposure without settlement, and liquidity without sovereignty. The asset can moon while the monetary thesis dies.

Wrapper Dominance Ratio (WDR)

Wrapper dominance has a stock component and a flow component, and they must not be divided into one another: exposure is measured *at* an instant, usage *over* an interval, so a single ratio would scale with the arbitrary length of the measurement window. WDR is therefore reported as a pair.

Stock form — what share of economic exposure sits in custodial or synthetic form, where E_t^{wrapper} and E_t^{native} are exposures at time t:

$$\text{CustodialExposureShare}_t = \frac{E_t^{\text{wrapper}}}{E_t^{\text{wrapper}} + E_t^{\text{native}}}$$

Flow form — how activity divides over a stated horizon h, where $V_{t-h:t}$ is volume across that window:

$$\text{WrapperActivityRatio}_{t,h} = \frac{V_{t-h:t}^{\text{wrapper}}}{V_{t-h:t}^{\text{native}}}$$

The stock form is bounded in $[0, 1]$. The flow form is not: it divides one volume by another, so it is bounded below by zero and unbounded above, and it is undefined when native volume is zero. Both are window-explicit. A rising stock share with a rising activity ratio indicates the asset may be financializing faster than it is becoming money. The horizon h must always be published alongside the flow form.

Inputs to track: ETF/wrapper AUM; custodied balances; exchange balances; wrapped token supply; native fee share; non-custodial settlement volume; shielded/private settlement volume; staking/collateral lockups; fee-burn coverage.

Why wrappers come first, and why that is not reassuring. The observation that first-generation use of a new medium wears the old medium's form [McLuhan 1964] — early print looked like manuscripts, early television like radio with pictures — supplies the intellectual lineage of this ratio, and it changes what a rising WDR should be read as. It is not, on its own, evidence of pathology: during a medium transition, consumption of the new asset through the old form is the expected initial condition, not an anomaly. What it means is that the WDR is measuring a *phase*, and that the thesis's claim is specifically about whether the transition completes. Wrapper dominance in year three is the manuscript stage of the bearer asset. Wrapper dominance in year twenty is a medium that never finished forming. The failure gate above exists precisely because the frame does not promise completion: first-generation use is always the old content, and the red lines, not the analogy, decide when waiting has become failing.

Gold ran this experiment already. The financial lineage is a century old, and it is worth naming because it removes any temptation to treat wrapper dominance as a novel crypto pathology. Paper claims on gold grew until they dwarfed the metal: unallocated bullion accounts, gold certificates, and eventually exchange-traded products came to represent exposure-to-gold on a scale that holding-of-gold never matched — the vault holds a fraction of what the paper promises, and the fraction is the finding. A shareholder in a gold ETF owns a claim on a trust that owns gold; a holder of unallocated bullion owns, in the stress-relevant sense, a deposit at whoever holds it. Both are exposure. Neither is a coin in a vault, and on the day it matters — the redemption run, the allocation freeze, the jurisdictional block — the difference between allocated and unallocated is the entire difference. Paper gold did not destroy gold. It demonstrated the exact pattern this ratio watches for: the financial claim outgrowing the bearer object it references, prospering for decades, and leaving the underlying's monetary character unchanged while the exposure stack did the volume. That is the precedent the WDR tracks, and the reason it is a stock *and* a flow rather than a single number.

Failure gate: If WDR rises for multiple quarters while native fee share, private settlement, and collateral lockups stagnate, the asset may remain investable but its **SoV-as-protocol** thesis

is weakening. This gate is added to the Red Lines in Section 30.6 and tracked on a dedicated Wrapper Dominance Board (Section 26.9).

13.7 The Market Realization Plane

The Wrapper Dominance Ratio is a **pair** — a stock share and a flow ratio — reported side by side and never divided into one another. The stock component measures how much ownership has migrated into custodial and synthetic form; it does not explain the *flow* mechanics by which those wrappers set prices. That requires one more construct.

Market realization is deliberately **not** an eighth layer of the stack. Layers 0–6 produce privacy, proofs, compute, settlement, and governance. Market realization produces none of those; it is the conventional financial machinery that represents claims on the resulting monetary object. It is orthogonal to the stack, so it is modelled as a *plane* around the stack rather than a layer within it.

Market Realization Plane

The set of external institutions and instruments through which claims on the native monetary object are represented, financed, allocated, and priced: exchanges and custody, spot ETFs, corporate treasury vehicles, index products, options and futures, leveraged and inverse ETPs, dealer swaps, prime-broker financing, passive mandates, systematic trend strategies, and rules-based or agentic treasury systems.

The result is two nested loops, and the thesis needs both to be instrumented.

The inner protocol loop determines whether the stack works:

$$\text{Create/Compute} \rightarrow \text{Prove} \rightarrow \text{Settle} \rightarrow \text{Verify}$$

The outer market-realization loop determines how the stack is financially represented:

$$\text{Narrative} \rightarrow \text{Wrapper} \rightarrow \text{Allocate/Lever} \rightarrow \text{Dealer Hedge} \rightarrow \text{Price} \rightarrow \text{Narrative}$$

The inner loop determines native monetary function. The outer loop determines market-price realization. Either loop can strengthen while the other weakens, and the outer loop can run for years on narrative alone.

The plane is not a precaution invented for a young asset. The reference asset of the entire financial system already has one, unacknowledged. When a central bank reads the long end as a message, the reply from market structure is that no unfiltered message exists for the Treasury curve. A measurable share of its demand is generated by benchmark rules and accommodated through leveraged intermediation (Section 3.1.3). The decomposition that would have to come

first—net demand by tenor, in risk units, separated by buyer and by decision rule, with a residual the authors are made to name—is the object this section specifies for the base asset. The difference is only that here it is built in from the start, where for the Treasury curve it would have to be reconstructed after the fact. If the price of the world's risk-free asset cannot be read without first identifying who is speaking and under what rule, the price of a monetary candidate certainly cannot, and that is the whole of Premise 9.

13.7.1 A Three-Stage Model

The combined thesis can now be stated as three linked but separable processes.

Stage A — regime pressure creates structural need. Let R_t index financial repression, surveillance, synthetic media, and AI concentration. Compute demand is deliberately *not* a component of R_t: it is a demand driver in its own right, entering alongside regime pressure rather than as part of the pressure index, so that $D_t^{\text{triad}} = g(R_t)$ is not partially defined in terms of its own output. Regime pressure generates demand for triad capacity:

$$D_t^{\text{triad}} = g(R_t)$$

This is the causal layer developed in Part I.

Stage B — protocol design converts need into native value. Let $\phi_t \in [0, 1]$ measure the effectiveness of the monetary design: required fees, burn share, collateral, issuance discipline, and non-bypassability. Then native captured demand is

$$D_t^{\text{native}} = \phi_t \, D_t^{\text{triad}}$$

As $\phi_t \to 0$, the system can remain enormously useful while the monetary object fails. This is *not* the Value Capture Lemma restated as a coefficient: the lemma is an AND-gate on five conditions, while ϕ_t is an **empirical capture fraction** those conditions constrain. The lemma says partial capture is possible only when the gate is open; ϕ_t measures how much of the triad demand actually flows through once it is.

Stage C — market structure realizes native value as a price path. Observed price depends not only on native demand but on aggregate exposure flows and the liquidity available to absorb them:

$$\Delta \ln P_t = \theta_t + I(Q_t, \Lambda_t) + \eta_t$$

where θ_t is the non-flow fundamental or native monetary drift — defined to *exclude* flow effects, so that all flow impact enters once, through the $I(Q_t, \Lambda_t)$ term, rather than twice — Q_t is aggregate exposure demand (a flow over the interval $[t-1, t]$), Λ_t is the liquidity and market-

impact state, and η_t is residual. Aggregate demand decomposes by source:

$$Q_t = Q_t^{\text{native}} + Q_t^{\text{allocation}} + Q_t^{\text{levered}} + Q_t^{\text{dealer}} + Q_t^{\text{trend}}$$

Only the first term is monetary evidence. The remaining four can dominate the price for extended periods, in either direction. Appendix H develops the flow mechanics formally — holder flow elasticity, the net mechanical gain coefficient, the recycling boundary, market impact, and volatility drag — and Section 26.8 defines the telemetry that makes the plane observable.

Why a protocol should care about external market structure. A designer might object that this is someone else's problem. It is not. External wrappers affect treasury behavior, collateral demand, governance concentration, liquidity-provider incentives, public understanding, security-budget expectations, regulatory pressure, and the price that the protocol's own participants use when making decisions. The protocol should not manage its price. It must nonetheless instrument the structures through which its monetary claims are represented, for the same reason it publishes VerifyPrice: unmeasured dependencies are where theses go to die quietly.

13.7.2 The Numeraire Is Not Fixed

Everything above measures the price path in units of account whose own value is assumed constant. That assumption is doing more work than it appears to.

Corollary 13.1 establishes that price is not evidence of monetary adoption. The point here is adjacent and independent: even a price series that *is* informative about the numerator can mislead through the denominator. A rising quotation may report an appreciating asset, a depreciating unit of account, or any combination of the two, and the series alone cannot distinguish them.

Numeraire-Dependence

The property that a measured price path is jointly determined by the asset and by the unit in which it is quoted, such that the same series can support opposite conclusions about real command over resources depending on the denominator chosen.

For this thesis the consequence is a measurement requirement rather than a philosophical observation. The Market Realization Plane exists to prevent financialization from being narrated as adoption. A plane instrumented only in the sovereign unit of account is blind to the case where wrapper-led appreciation and currency depreciation move together—which is precisely the regime Section 3.1 argues is likely.

Requirement. The core market-realization series—the Wrapper Dominance Ratio (Section 13.6), the Wrapper–Native Growth Gap, and the native fee, burn, collateral, and settlement series—should be published in at least one non-fiat numeraire alongside the fiat series. A reference basket of energy, compute, shelter, and gold is sufficient and has the advantage that each component is a thing the holder might actually need. Where the two numeraires disagree in sign, the disagreement is the finding and should be surfaced rather than reconciled.

A sharper definition of the objective. The thesis has so far spoken of purchasing power, which inherits the numeraire problem it is trying to escape. The more precise statement of what a store of value is for:

> Wealth is durable, transferable command over necessary capacity under adversarial conditions.

Each term is load-bearing. *Durable* excludes claims that survive only while an issuer chooses to honor them. *Transferable* excludes capacity that cannot be moved or bequeathed. *Necessary capacity* names the denominator explicitly—energy, shelter, compute, mobility, time—rather than leaving it implicit in a currency. *Under adversarial conditions* is the requirement of Chapter 4 and the reason Item 9 is a design constraint rather than a preference. This is the definition the rest of the document should be read against.

Worked example: structure does not determine behavior. A useful illustration, because it is easy to get backwards. An asset can be genuinely bearer-oriented, scarce, censorship-resistant, and independently settleable at the protocol layer, and simultaneously trade as leveraged technology beta at the wrapper layer—held through exchange-traded products, pledged as collateral, sized by volatility-targeting mandates, and sold first in a de-grossing episode. Both descriptions are accurate. They describe different planes.

The error in either direction is the same error. Concluding from correlated selloffs that the protocol properties are illusory mistakes the plane for the stack. Concluding from the protocol properties that the price will behave defensively mistakes the stack for the plane. Appendix H supplies the machinery for keeping them apart, and Section 26.8 the telemetry.

A related precision, since the language is routinely abused: a proof-of-work asset is energy-*linked*, not energy-*backed*. Its issuance and security consume energy; it is not redeemable for a fixed quantity of electricity, and no holder has a claim on a joule. The distinction matters because a genuine claim on capacity—a generation asset, a storage system, a Facility Capacity Receipt—has a payoff structure that an energy-linked bearer asset does not, and conflating them overstates the latter's hedging properties.

Technology success is not incumbent success. One further reading of the same distinction, which the thesis has stated elsewhere in monetary terms and which generalizes. A technology can transform an economy while destroying the equity, debt, and tokens of everyone who financed its first build-out; railways and long-haul fiber are the standard cases. This is Corollary 13.1 pointed at infrastructure rather than at protocols: the social return and the return to the first capital structure are different quantities, and the second is the one an investor actually receives.

13.7.3 Denomination: The Function the Thesis Has Never Measured

Ask a merchant in Buenos Aires what a thing costs and the answer comes in dollars, whatever ends up crossing the counter. That is the whole of the unit-of-account question, and it is the one monetary function the other instruments do not measure.

Section 13.7.2 treats the numeraire as a measurement hazard: the thesis's own series are quoted in a unit whose value moves, so they must also be published in a non-fiat one. This subsection points the same word the other way, and asks a monetary question rather than a hygiene one — *is anyone choosing to quote in this one?*

Denomination is the strongest of the monetary functions and the one nobody instruments. A store of value can be abandoned in an afternoon: the holder sells and is out. A unit of account cannot. It is welded into contracts whose counterparties would all have to move at once, and each of whom eats the basis if they move alone. That mutual dependence is what makes denomination sticky: it is a coordination device rather than a portfolio choice, and coordination devices are hard to acquire and correspondingly hard to lose [Doepke & Schneider 2017]. Assets that achieve denomination rarely lose it. Assets that achieve only holding lose it routinely.

The thesis measures fees, burns, collateral, settlement, holder cohorts, wrapper exposure, capacity, reachability, and verification cost. It does not measure whether a single contract anywhere is *priced* in the unit. That is the largest gap in a five-family measurement programme, and it is a gap on the monetary side rather than the service side.

The trap, stated first, because the obvious metric is worthless. A naive denomination share would read one hundred per cent on the day the network launches and would mean nothing. Section 8.4 specifies that all triad services are priced in the base token; Condition 1 of the Value Capture Lemma requires it. Denomination that the protocol *mandates* is a routing rule — the same routing rule already stated, already conceded to be a design commitment rather than a finding, and already excluded from carrying monetary weight by Section 13.3.9. Counting it as evidence of adoption would be the tokenomics inference in a new costume.

Everything of interest is therefore in the complement.

Free-Choice Denomination Share

Let $\mathcal{C}_{\text{free}}$ be the set of contracts in which the parties were *not* required by protocol rule to express price terms in the base asset — forward and term capacity purchases, SLA agreements above the mandated fee, bilateral prover–buyer contracts, corridor quotes, over-the-counter blocks, service retainers, employment and vendor agreements among operators, and any obligation denominated between counterparties rather than by the protocol.

Free-choice denomination share is the fraction of $\mathcal{C}_{\text{free}}$ whose *price terms* are written in the base asset, reported by contract count, by notional, by tenor bucket, and by counterparty class.

The mandated complement — fees, required collateral postings, protocol-set tier pricing — is published separately and is **never** added to it. It is a design commitment and reads as one.

Quoting and settling are different acts, and the crypto record turns on the difference. A contract has two units and they need not agree. It can be quoted in dollars and settled in the asset; quoted in the asset and settled in the asset; or, rarely, quoted in the asset and settled in dollars. The second and third are unit-of-account observations, because the unit that fixes the obligation is the unit of account whatever instrument discharges it. The first — **quoted in fiat, settled in the asset** — is a payment-rail observation, and it is the configuration essentially all crypto commerce has actually taken.

This is worth stating at full strength because it is the strongest evidence against this thesis's own expectations. **Bitcoin is the most successful new store of value of the last century on the measures this thesis uses—credible scarcity, verification cost, censorship resistance, and capitalization—and after seventeen years it denominates almost nothing.** Bitcoin-native businesses quote in dollars and settle in bitcoin. Salaries described as paid in bitcoin are, with few exceptions, dollar salaries settled in bitcoin. Whatever monetary premium bitcoin carries, it did not come with denomination attached and has not produced it since.

Two readings are available and the thesis should not pretend the choice is obvious. Either denomination is *not* necessary for moneyness — in which case this subsection measures something real but not load-bearing — or bitcoin's moneyness is incomplete along a specific, measurable axis. This document takes the second reading and concedes it is contestable.

The one structural asymmetry that makes the second reading defensible. Bitcoin has no native economy of goods. The things bought with it are produced by an economy that prices in dollars, so a bitcoin quote imports basis risk against every input the seller must pay for, and no coordination is available to remove it. A triad protocol is different in one respect that is not

a matter of degree: *the goods are produced by the system that issues the unit.* A proof, a verified inference, a corridor settlement, and a unit of delivered verified capacity are priced, produced, collateralized, and settled inside the same accounting perimeter. A prover paid in the unit, posting collateral in the unit, buying capacity priced in the unit, has a naturally two-sided position and faces the basis risk that keeps bitcoin quotes in dollars only on the legs that reach outside.

That is the whole of the structural argument, and it is narrow. It says a service protocol has a coordination path to denomination that a pure bearer asset does not. It does not say the path will be taken.

The walled-garden test, before a critic applies it. Internal denomination proves nothing on its own, and there are unflattering precedents. Airline miles are denominated, transferable, and priced against a schedule the issuer sets; in-game currencies denominate large and genuinely productive internal economies. None is money, and each would score highly on a share computed only over contracts internal to its own economy. The discriminating cut is therefore not free choice alone but *where the counterparty stands*:

- **Tier I — inside contracts.** Both legs are protocol services. Free choice, but the parties are already inside the accounting perimeter. This is where in-game currencies score one hundred per cent.
- **Tier II — boundary contracts.** One leg is a protocol service, the other is not: an enterprise buying verified inference, an operator's power contract, a hardware purchase.
- **Tier III — outside contracts.** Neither leg is a protocol service, and the unit was chosen anyway.

The monetary content rises strictly from I to III, the three are published separately and never blended, and **Tier I is reported as context rather than as evidence.** A protocol scoring highly on Tier I alone has built a company scrip with excellent telemetry.

The volatility objection, which is the serious one. Nothing prices in a volatile unit if it can avoid it, and this is the standard and correct explanation for the bitcoin record. It bounds where denomination can appear, and the bound is useful precisely because it makes the prediction specific rather than directional. Denomination should appear **first at short tenor**, where the parties' exposure to unit variance over the contract's life is small, and **first in two-sided positions**, where a counterparty's costs and revenues are in the same unit and the variance nets. It should appear **last, or never**, in long-dated one-sided contracts, which is exactly where a naive story would expect a strong monetary asset to show up first.

That is a shape prediction across the tenor and sidedness buckets, and a shape is harder to fake than a level. A denomination share that rose uniformly across all buckets would be evidence of subsidy or mandate rather than adoption, and should be read that way.

And it couples to the previous mechanism. The reason volatility is prohibitive is not variance itself but the cost of carrying and hedging the exposure that variance creates over a contract's life. That cost falls when positions in the unit can be financed cheaply and pledged without investigation — which is what Section 13.3.10 argues cheap universal verification supplies. The two mechanisms are not parallel additions to the same list. **Information-insensitivity lowers the carrying cost of denomination**, and denomination is what converts an accepted asset into a coordination device. Acceptance without denomination is a settlement rail; denomination without acceptance does not occur.

An ordering caveat the thesis should carry rather than discover. The currency-substitution literature finds a consistent sequence when populations abandon a domestic unit: asset substitution first, unit of account second, means of payment last. Savings dollarize years before prices do. If that ordering holds here, free-choice denomination is a *confirming* indicator of monetary adoption rather than a leading one on the macro path, and this subsection should not be read as promising an early warning of anything.

The claim it does support is narrower and survives the caveat: denomination is *leading with respect to the protocol's own contract layer*, because Tier I and Tier II contracts exist from Phase I onward and can be counted long before any external population is substituting anything. What is early here is the measurement, not the phenomenon.

Instrument. Publish quarterly on the Value Capture Board (Section 26.9), under the custody and reproducibility terms of Red Line 14 — the party collecting the fee may not classify the contracts:

- **Free-choice denomination share**, by contract count and by notional, split across Tiers I/I-I/III, and *never* aggregated with the mandated complement.
- **Quote/settle matrix.** The four-cell table of quoting unit against settling unit, by notional. The fiat-quoted, natively-settled cell is the payment-rail null and is reported as the headline denominator, not hidden in a total.
- **Tenor and sidedness decomposition.** Share by tenor bucket and by whether the counterparty's position is one- or two-sided in the unit, so the shape prediction above can fail visibly.
- **Subsidy disclosure.** Any fee rebate, tier privilege, or incentive conditioned on native denomination, with the share recomputed excluding subsidized contracts. A purchased quote is not a chosen one.
- **Redenomination events.** Contracts that switch units in either direction at renewal, which is the cleanest available revealed preference and the series a critic will ask for first.

The falsifier. This mechanism is made falsifiable as **Red Line 18** (Section 30.6), in two conditions with different severities: outsiders declining to reckon in the unit weakens the

thesis, and protocol participants declining to reckon in it on unmandated contracts kills it.

What this subsection does not claim. It does not claim denomination is necessary for moneyness — bitcoin is the standing counterexample and it is named above rather than managed. It does not claim the protocol's endogenous service economy makes denomination likely, only that it makes a coordination path available that a pure bearer asset lacks. It does not raise any rating in Table 8.1: a new instrument on link 9 is an instrument, not a finding. And it does not promise an early warning on the macro path, for the ordering reason stated above.

What it does supply is an answer to the question the other instruments cannot answer: *what would be visible, at Phase I, that distinguishes this from every service token that ever shipped?* The answer is a contract nobody was required to write in the unit, written in the unit anyway, with a counterparty outside the perimeter, at a tenor long enough to matter, without a rebate.

Chapter 14

Service-Instrument Design Space

Once work can be verified, the question is what to build on top of it. The triad produces three capacities, and Work Credits turn capacity into claims. But a claim is a design decision, not an inevitability, and the decision has consequences that outlive whoever made it. Pick the wrong instrument for a real buyer, and the buyer quietly routes around you: a treasury that needs a two-year hedge will not buy a perpetual index, and an AI lab hedging inference cost will not buy privacy credits it cannot redeem. The failure mode is not loud. Nobody announces they are leaving; the service just never gets the volume, and the instruments sit on the shelf with impeccable specifications and no counterparty.

That is why this chapter walks the design space rather than picking from it. Monetary candidacy remains confined to the base asset and subject to the Triad Coherence Test (Section 8.3); what follows is the inventory of everything else you can build, organized by which triad leg the claim draws on, and by who — concretely — would ever hold it.

The legacy labels are kept because they name real demand clusters, but read them as *lenses*, not products. "ZK Money" is not a coin; it is the observation that privacy and proofs are usually bought together. "AI Money" is not a token; it is the observation that compute is bought forward, in size, by parties who fear shortage. Each lens describes a buyer, and each buyer implies a different instrument.

14.1 The "ZK Money" Analytical Lens

The legacy "ZK Money" label groups service instruments that primarily reference **Privacy + Proofs**. In practice this means claims on shielded settlement capacity — corridor LP positions, privacy-rail credits — and claims on zero-knowledge proof capacity such as SNARK/STARK service tiers.

What these share is a particular pairing: bearer-like, censorship-resistant settlement rights, bundled with attestations of compliant behavior. Neither half is worth much alone. Settlement rights without attestations are an offshore account with better cryptography; attestations without settlement rights are an audit trail nobody can act on. The value is in the bundle, which is why regulated actors — who need privacy from counterparties and provability toward supervisors *simultaneously* — are the natural buyer.

One transaction to fix the idea. A corporate treasurer operating under an emergent capital-control regime needs to move payroll to overseas staff without exposing vendor rates to competitors or the transaction graph to the local banking authority, and must still satisfy the parent company's auditors at quarter-end. She buys corridor credits that settle shielded, and each settlement emits a PIDL receipt her auditors can verify without a viewing key. That single purchase — privacy toward the state, disclosure toward the auditor, both on purpose — is what "ZK Money" was always pointing at. Treasuries under repression, regulated actors needing simultaneous privacy and compliance, and individuals wanting cash-like digital assets are variations on the same shape.

14.2 The "AI Money" Analytical Lens

The legacy "AI Money" label groups service instruments that primarily reference **Compute + Proofs**: claims on verified FLOPs, futures on model-specific inference capacity with proof guarantees, and equity-like positions in proof factories whose output is standardized, verified compute.

The common thread is that compute is the one triad leg bought *forward*. Nobody pre-pays for privacy next year; plenty of parties pre-pay for inference next year. That asymmetry is the entire commercial logic of the lens: it exists because shortage risk, not usage, is what generates a term structure.

One transaction. An AI lab training a frontier model has a contractual delivery window for a customer deployment and has watched spot GPU prices move 40% in a quarter. It buys futures on verified inference for the specific model hash it will serve, with proofs of correct execution attached, locking both price and correctness for the delivery window. The hedge is only as good as the verification — an unverified forward on compute is just a promise from whoever sold it — which is why the claims are typed to workload and proof tier rather than to "compute" in the abstract. AI labs and application builders hedging compute shortage, and investors who believe AI demand outlives any particular model cycle, are the same buyer at different horizons.

14.3 Hybrid Instruments

Hybrids reference all three legs at once, and they exist because real obligations rarely arrive one leg at a time.

A **triad basket** is an index-like instrument backed by diversified Work Credits across privacy, proof, and compute workloads — for an allocator who wants exposure to the stack's aggregate throughput without forming a view on which leg tightens first. A **corridor-linked compute claim** couples private settlement rights with AI inference capacity: the holder can

pay for and receive verified inference without revealing who is asking, which is the shape an agent acting for a pseudonymous principal actually needs. An **agent service treasury** takes the same idea and makes the holder the agent — an on-chain system holding combinations of typed proof, settlement, and compute claims to fund its own operation, buying its own compute the way a factory buys its own electricity.

One transaction, and the one worth watching. A software agent under contract to serve a customer's inference workload for a quarter holds its own treasury of corridor-linked compute claims, draws them down as it serves, and settles with its customer over privacy rails — no human in the loop, and every drawdown receipted. If that transaction ever runs at scale, the instruments holding it up are hybrids by construction. Chapter 7 models the buyer this transaction implies, and Section 7.5 states why an agent treasury must be classified by principal before it counts as a holder.

The key design constraint is always the same, and it is the chapter's answer to the question it opened with: **claims must stay tightly coupled to verifiable work and capacity**, not to governance mood or marketing narrative. The moment a claim's value can be moved by a vote or a slogan, it has stopped being a service instrument and started being a governance lottery — and the buyer who needed the hedge is already gone.

Chapter 15

The ZK + AI Service Economy

ZK is the accounting engine of this stack. If privacy is the right to speak softly and compute is the ability to think loudly, zero-knowledge proofs are the receipts that make both negotiable. They turn "trust me" into "verify me" without exposing the underlying state. A world that routes value and decisions through proofs rather than paperwork is, in effect, a **zk-economy**.

A zk-economy is not just "more proofs." It is day-to-day life running on attestations: every payment, inference, bridge, and audit backed by succinct evidence that any party can check.

15.1 Why Open Hardware Is a Precondition for the ZK-Economy

A sound proof does not need an honest prover. That is what soundness means: a machine that lies about the computation produces a proof that fails to verify, whoever built the machine. What a dishonest machine can still do is everything soundness does not cover.

Today's largest ZK systems run on stacks that are almost entirely closed: proprietary GPUs, opaque microcode, black-box TEEs, firmware that can be updated silently overnight. A government-mandated "secure enclave" with a secret backdoor can exfiltrate witnesses from supposedly private circuits; the proof verifies and the privacy is gone. A trusted setup whose toxic waste survived on someone's hardware lets its holder forge proofs outright—a failure of the ceremony rather than of the prover, but one that lives on machines all the same. And every claim that reaches past the circuit—that this camera captured this scene, that this plant burned this energy, that this prover won this slot fairly—is only as good as the device that measured it.

Closed hardware leaves math-grade correctness wrapped around vendor-grade trust.

A genuine zk-economy therefore begins one layer below the circuits, with **verifiable machines** (Layer 0). At least part of the proving path must be open from the RTL up through the software stack.

Perfect certainty about every chip in circulation is not available, but the stack can move to a regime where:

> *"This proof came from this class of machine, built from this design, and it would have been extraordinarily expensive to tamper with it without being caught."*

That is enough to make "trust the hardware" a falsifiable claim instead of a sacrament.

Concretely, a zk-economy that wants to support the service markets described by the "ZK Money" and "AI Money" lenses needs three things from Layer 0:

- **Open proving paths:** For canonical circuits, at least one proving stack must be open from RTL through firmware and prover binaries.
- **Sampled, attestable devices:** Lots can be sampled with verifiable randomness, imaged, and subjected to structured tests.
- **Hardware-level SLOs for witness confidentiality and physical claims:** Profiles with weak sampling or poor audit history should be priced differently.

15.2 How This Opens a ZK + AI Economy

The same logic extends to AI. If proofs are the receipts of the digital order, AI is the industrial plant that consumes energy and data and emits capabilities.

The zk + AI economy is built on three interacting strata:

1. **AI blockchains ("model chains"):** Duplex-style and Ambient-style [Ambient 2024] PoUW systems treat model training and inference as work functions. Work Credits sold against these workloads remain typed verified-compute service claims. Pass [Pass 2026] formalizes the regime in which token rewards subsidize inference below compute cost ("Duplexia"); the fee-incidence bound of Section 13.3 says such a subsidy cannot be funded from fees and is therefore paid from issuance, which is the design's cost of capital, not evidence of demand.
2. **ZK blockchains ("proof chains"):** Nockchain-style [Nockchain 2025] zk-PoW systems coordinate global prover markets and issue typed claims on future proof capacity.
3. **Privacy rails as connective tissue:** BTC↔XMR swaps (BTC↔shielded-ZEC gated on the work in Section 23.3), shielded pools, and private rollups let capital move without custody.

The daily-life version looks almost boring:

- Your phone proves you paid a toll **without revealing your route**.
- A DAO buys inference on a model chain; a proof chain clears the resulting proofs; payment moves over privacy rails.
- An exchange settles cross-chain obligations with batches of zk-proof-backed settlement receipts.
- An insurer prices climate risk from sensor networks whose readings are baked into proofs from open hardware profiles.

Underneath the surface, three demand curves reinforce one another:

1. **Proof demand grows** as systems move from "trusted unless flagged" to "untrusted unless proven."
2. **Verified compute demand grows** as AI saturates workflows.
3. **Privacy demand grows** as financial repression and surveillance tighten.

These are three economic flywheels that, if the thesis is right, will spin up over time:

Proof flywheel. More systems demand proofs → prover markets deepen → VerifyPrice falls or stabilizes → more systems can afford to demand proofs.

Compute flywheel. More AI in workflows → more willingness to pay for verified compute → more Work Credits sold against AI capacity → more capital available to fund model chains.

Privacy flywheel. More repression and surveillance → more demand for lawful private rails → deeper anonymity sets and corridor liquidity → cheaper and safer private settlement → more users adopt privacy by default.

Open hardware is what keeps these flywheels from collapsing into a handful of "trust me" platforms. zk-PoW is what coordinates proof work into a measurable, meterable commodity. AI-PoUW is what turns model work into an asset class whose revenues are natively tied to **triad usage**. Privacy rails are what keep capital flowing between them without recentering custody.

Put in one line:

> **ZK makes computation checkable by anyone. Open hardware keeps the witness private and the physical claims honest. AI gives us something valuable to spend that checking on. zk-PoW and AI-PoUW are the mechanisms that weld the three into an economy rather than a collection of clever demos.**

Chapter 16

SoV Evaluation Framework

Running this framework is not scoring a product. It is auditing a supply chain that runs in two directions at once — upward from machines into money, and outward from one balance sheet to whoever is standing behind it. The operator/investor checklist in Part V spans the whole stack; this chapter is the **monetary lens**, which evaluates whether the base asset is a credible store-of-value candidate while keeping the service and evidence instruments beneath it analytically separate.

Every question below is a way of asking one of two things: *can this thing deliver what it promises under the conditions that make it valuable?* and *who absorbs the loss when it cannot?* The first is the question a buyer asks; the second is the question an underwriter asks. A store of value has to survive both, which is why the checklist reads as though written by two people.

One distinction does more work than any other here, so it belongs at the front rather than buried as a note: **legibility is not resilience**. A facility can be perfectly transparent about being catastrophically fragile — publish every receipt, pass every audit, display its full telemetry, and still fail on the first stressed day. The reverse is rarer but real: a robust operation whose interior is opaque. Every measurement regime, this thesis's included, is biased toward the first thing it can see; dashboards reward the visible. So the checklist deliberately separates the questions that ask *can the substrate be seen* from the ones that ask *can what we see withstand anything*, and treats a high score on the first set as establishing nothing at all about the second. Items 10 and 11 exist as two items for exactly this reason. Passing the legibility questions makes an asset *auditable*; only the resilience questions can make it a store of value.

For a candidate base asset X, ask:

1. **What triad services does it route value from?**
   - Exactly which Privacy, Proof, and Compute flows pay fees, burn supply, or require base-asset collateral?
   - Are those workloads canonical and well-specified?
   - Are the associated Work Credits reported as typed service claims rather than counted as monetary holdings?
2. **How is base-asset issuance constrained?**
   - Is issuance formulaic and transparent, and distinct from Work Credit issuance against DVC?

- Can governance mint outside of those rules? Under what constraints?

3. **Can anyone verify the service linkage?**
   - Are VerifyPrice and workload telemetry public?
   - Can a third party recompute DVC and sample the proofs, settlements, and Work Credit claims that justify fee and collateral activity?
4. **How does it behave under repression?**
   - If real yields are -300 bps for 3 years, what happens to demand for X?
   - If on-/off-ramps are throttled, can X still be acquired, held, and spent?
5. **Is privacy lawful by design?**
   - Are there clean paths for voluntary disclosure that do not reintroduce custody or KYC chokepoints?
   - Are the receipts (PIDL) expressive enough for auditors?
6. **What is the duration profile?**
   - Does X promise fixed nominal coupons (bond-like) or variable participation in fee flows (equity/commodity-like)?
   - Can a sovereign push its real return negative by fiat?
   - If X is duration-neutral as money, what *separate* instrument, if any, is supposed to warehouse the construction interval of the physical assets the stack requires? Conflating the two is a completeness failure, not a feature (Section 3.3, Section 34.5.1).
7. **Who can be locked out?**
   - Is access to X gated by one jurisdiction, cloud, or app store?
   - Do Layer-0/1/2 assumptions create hidden chokepoints?
8. **Can users bypass the asset?**
   - Can users buy equivalent Privacy, Proofs, or Compute service for fiat, stablecoins, or direct cloud contracts without touching the asset?
   - Are fees actually paid in the native asset, or do operators accept and immediately off-ramp alternatives?
   - Are burns and collateral lockups material, or negligible relative to total supply?
9. **Can institutions legally and operationally hold it?**
   - Is custody support available from regulated custodians?
   - Is the legal classification clear enough for treasury or fund mandates?
   - Is liquidity deep enough for institutional-scale positions?
   - Does governance redirect value away from holders?
10. **Is the physical infrastructure behind it legible?**
    - Are Facility Capacity Receipts (Section 17.7) available and independently verifiable for a material share of active capacity?
    - Is Physical VerifyPrice (Section 22.8) within published bounds, or does verifying the facility take months of trust-the-vendor diligence?

- Is supply concentrated in a small number of hyperscalers, closed TEEs, or single-jurisdiction facilities (Homestead Ratio / Enclosure Risk Flag, Section 5.6)?

11. **Is that infrastructure resilient, or merely visible?**
    - What is network sovereign optionality, $\mathcal{O}_s^{\text{net}}$ (Section 17.8), and is it published with its component vector and weights rather than as a bare scalar?
    - What share of VERIFYPRICE-tracked capacity sits in jurisdictions with an active curtailment, rationing, or interconnection-denial regime (Section 5.5)? Is Red Line 13 approaching?
    - Is the headline verification cost a benign-state figure, or is disruption-adjusted VerifyPrice published alongside it with stated disruption probabilities (Section 17.10.2)?
    - Does the resilience premium demonstrably lower disruption-adjusted cost, or is it capital misallocation with a security narrative attached (Section 17.10.3)?
    - Is issuance bounded by stressed Delivered Verified Capacity, with the active binding constraints named, or by nameplate power and installed hardware (Section 17.9)?

    This is the pair to the previous question, and an operator will try to collapse them. The prior question asks whether the substrate can be *seen*; this one asks whether what we see can *withstand* anything — the distinction opened at the head of this chapter, and the reason Red Lines 10 and 13 are separate.
12. **Is demand wrapped, or protocol-native?**
    - What share of economic exposure sits in ETFs, treasury companies, and custodial wrappers versus protocol-native settlement (Wrapper Dominance Ratio, Section 13.6)?
    - Is price rising while native fee share, private settlement volume, and collateral lockups stagnate?
13. **Does anyone hold it through stress?**
    - Is there a persistent, self-custodied, loss-bearing cohort that adds through drawdowns (net of units it spends on the service during stress, which count for it), or is native demand acquired just in time, sold on receipt, leveraged, or wrapped (Native Monetary Buyer Map, Section 26.8.8)?
    - Does holding in advance beat Bitcoin plus conversion on the day, for any workload and stress state the access experiment has run (Section 13.1)?
    - Is the persistent bid a rule that absorbs less risk as prices fall, or leverage that leaves on a margin call (Section 3.1.3, Section 3.1.3)? A large bid of either kind is not an anchor.
    - Where holders are agents, who is the principal, and can a person still revoke the agent through the protocol? Does the cohort move as one policy (Section 26.8.9)?
14. **Is anything monetary happening, or only accrual?**
    - *Pledgeability.* Do lenders disagree less about haircuts on verifiable positions than on opaque comparables of the same volatility, and re-converge faster after a shock (Section 13.3.10; Red Line 17)?

- *Denomination.* Does anyone write contracts in the unit that nobody required them to write in it, and do the protocol's own participants denominate their short-tenor obligations in it (Section 13.7.3; Red Line 18)?
- *Holder-side service.* Does the price of borrowing or parting with the asset rise in the pressure corridor above the pre-registered minimum increment (Section 13.3.11; Red Line 16), and does a held balance beat Bitcoin converted on the day (Section 13.1)?

This is the only question on the list that reads on moneyness itself. Every other question asks whether the conditions for it hold. Fees, burns, and collateral cannot answer it, because what they establish is a cash-flow claim and a collateral lockup (Section 13.3).

15. **Does it equip holders, or manage them?**
    - What share of flows are non-custodial and require no global identity (Item 9)?
    - Does disclosure stay selective (predicates) rather than converging into a reconstructable dossier (Section 27.1.2)?

A base asset that scores well on these questions remains a **next-generation SoV candidate**:

- It is backed by verifiable work, not by decree.
- Its value is rooted in capacities that a digital civilization must keep buying: privacy, proofs, and compute.
- It is measurable under stress via telemetry, not defended by rhetoric.

16.0.1 Comparison: Traditional Hard Assets vs. the Base-Asset Candidate

For context, it is useful to compare how gold and silver—the canonical "old-world" stores of value—score against the same framework:

Gold and silver excel on **Lindy** (millennia of history), **simplicity** (no software, no network), and **zero counterparty risk in physical form**. They remain rational holdings during the transition window while the triad stack matures. Healthy VerifyPrice, VerifyReach, VerifySettle, DVC, and native-buyer telemetry can establish distinctive digital capabilities; they do not establish that the base asset is superior money or promote Work Credits into savings instruments.

The thesis does not claim gold is obsolete. It claims only that a base asset may earn a distinct premium for preserving agency across specific digital failure modes—financial repression under surveillance, synthetic media, and compute enclosure—if all nine links hold. The bridge is still load-bearing; the candidate remains under test.

The rest of the thesis will shift back to the **stack angle** (Layers 0–6) and the **telemetry angle** (VerifyPrice/Reach/Settle, DVC, buyer quality, decentralization, governance). The monetary question stays in the background: every design choice is judged by whether the base asset can preserve agency and satisfy the nine-link chain, while Work Credits stay what Section 12.8 established them to be — service instruments — under every regime.

Criterion	Gold / Silver	Triad Base Asset
What supports it?	Geological scarcity; 5,000-year Lindy effect	Conditional holder-side agency plus non-bypassable value accrual from stress-deliverable triad services
Issuance discipline?	Mining (energy + geology)	Constitutional base-asset schedule; separate from DVC-bounded Work Credit issuance
Anyone can verify?	Requires assay or trusted custodian	Laptop-class verification in seconds
Behavior under repression?	Seizure risk at borders; detectable	Shielded settlement; non-custodial corridors
Privacy by design?	Physical possession is private; custody reintroduces counterparty	Default-encrypted; viewing keys for voluntary disclosure
Duration profile?	Zero yield unless lent (counterparty risk)	Duration-neutral base asset; project credit remains separate
Who can be locked out?	Jurisdiction can ban ownership, seize at borders	Layer-0/1/2 resilience; multi-path reachability
Programmability?	None	SLAs, escrow, conditional logic, composability
Settlement speed?	Days (physical); hours (paper with counterparty)	Minutes (privacy corridors)

Table 16.1: SoV comparison: traditional hard assets vs. the conditional triad base-asset candidate. Work Credits are excluded because they are typed service claims.

Part III
Infrastructure Layers 0–3

Part I argued that soft guarantees are breaking under debt, AI, and platform control, set out a threat model plus a seven-layer cypherpunk stack, and named the buyer: increasingly a machine delegating to another machine. Part II treated **Privacy, Proofs, and Compute** as distinct services that may support a conditional **base-asset monetary candidate**; Work Credits remain typed, energy- and capacity-anchored service claims rather than claims on a monetary base.

Part III switches to the **infrastructure angle**. It asks a blunt question:

> *What has to be physically true, and stay true under repression, for the triad to be real?*

The answer starts at **Layer 0 – Verifiable Machines & Energy**, then climbs through **Layer 1 – Reachability**, **Layer 2 – Distribution & Execution**, and **Layer 3 – Identity & Claims**. These layers do not directly "store value," but they determine whether any higher-layer SoV story survives contact with hardware, networks, and identity regimes.

Chapter 17

Layer 0: Verifiable Machines & Energy

Layer 0 is where cryptography stops being metaphor and touches matter.

PoW had an implicit Layer 0: silicon, power, and warehouses hashing in the dark. The trust assumption was: *"ASICs will do what the SHA-256 spec says."* In practice that meant: *"we trust the vendor, the fab, the firmware, and the power company, and we hope no one has a better ASIC they haven't told us about."*

Those assumptions were never harmless—a compromised fab or a secretly superior ASIC bears directly on consensus security, and every holder relies on that—but the damage was confined to one quantity, the distribution of hashpower, and the hash itself carried no information anyone outside the system consumed. They stop being tolerable in a new way when the puzzle mints **receipts** that other people act on, and when proofs and verified compute become monetary primitives: a compromised machine can then misreport what was computed, on which hardware, at what energy cost, and third parties who never touched the consensus layer act on the misreport. Giving the work an external buyer raises the Layer 0 requirement rather than lowering it.

17.1 Why Layer 0 Is a Monetary Question

Gold worked as money because geology is hard to fake at scale. Bitcoin worked because hashing cost was hard to fake at scale. In both cases, the monetary story rested on silent assumptions: rocks behave; fabs behave; physics behaves.

The triad inherits those assumptions and tightens them:

- Privacy is only as real as the devices that hold keys and speak on the wire.
- Proofs are only as real as the machines that generate entropy, execute circuits, and sign receipts.
- Compute is only as real as the GPUs/ASICs that claim to have run workloads.

If those machines are opaque, remotely steerable, or quietly biased, **Work Credits degrade into theater**:

- A compromised RNG can turn "unpredictable leader election" into a slow rug-pull.
- A backdoored prover can leak witnesses or mishandle private inputs, even though it cannot emit invalid proofs accepted by a sound verifier unless the proof system or implementation

is broken.

- A mandated TEE can become a kill switch for entire clusters of provers and routers.

All of this math still runs on matter.

Every proof, every encrypted wallet, every verified FLOP ultimately lives on a sliver of doped silicon that almost nobody is allowed to audit. Today's "trusted hardware" stack is a daisy chain of NDAs: closed-source EDA tools, proprietary IP blocks, opaque PDKs, black-box fabs, sealed packaging, vendor-run attestation services.

Layer 0 Precision

Layer 0 does not make cryptographic proofs sound; soundness comes from the proof system and verifier. Layer 0 makes claims about the physical world credible: which machine produced a receipt, whether inputs were captured honestly, whether randomness was biased, whether witnesses were protected, and whether energy/capacity claims are real. A sound proof system should not require trusting the prover's hardware for proof correctness. Layer 0 becomes essential when the claim includes physical capture, machine identity, energy use, witness confidentiality, randomness quality, side-channel resistance, or fair participation in useful-work markets.

From a monetary standpoint, Layer 0 asks:

> *"Can we treat triad capacity as collateral if we don't know what the machines are really doing?"*

The answer is "no." Triad services delivered by machines nobody can interrogate cannot support a credible base-asset monetary candidate; they are IOUs on a hardware cartel plus the jurisdictions that regulate it.

So the mandate of Layer 0 is:

> **Translate "trust the vendor" into "trust these verifiable claims about the machine and its power," or do not pretend it is money.**

17.2 Design Goals and Non-Goals

Layer 0 has to be ambitious enough to matter and humble enough not to LARP full supply-chain omniscience.

Goals

1. **Verifiability over purity.** The aim is *checkable claims* about machines, not metaphysical purity. Open RTL where possible; structured sampling where not.

2. **Common knowledge of security.** Different actors should be able to agree on *facts* about hardware profiles, even if they disagree about policy.
3. **Energy anchoring, not energy worship.** Power use should be measurable enough that "Work Credit per joule" is meaningful.
4. **Degradability under attack.** When assumptions fail, the system should *degrade visibly*: telemetry spikes, profiles are deprecated, Work Credits tied to broken profiles are risk-flagged and may face prospective ineligibility for new collateral uses.
5. **Composable exports.** Layer 0 should emit artifacts that higher layers can consume mechanically.

Non-goals

1. **Perfect trustlessness.** Layer 0 will not "solve" global hardware and supply chains.
2. **Single-vendor dependence.** Heterogeneity is a feature, not a bug.
3. **Total hardware transparency on day one.** Political and commercial realities exist.
4. **Magical protection against all side-channels.** The stack assigns **budgets** to attack surfaces.

17.3 Hardware as Base Reality for Work Credits

In Part II, Work Credits were defined as **claims on standardized units of triad work** (privacy settlement, proofs, verified compute) anchored to energy and VerifyPrice.

Layer 0 defines the **hardware profile** that each Work Credit type rests on. For a canonical workload W, a **hardware profile** H might specify:

- Chip family and stepping.
- Microarchitectural features (e.g., presence of certain accelerators or TEEs).
- RNG source and test regimen.
- Power metering and thermal envelope.
- Known limitations (e.g., "avoids TEE X due to backdoor Y; uses open core Z instead").

When a Work Credit of type (W, T) is fulfilled, its receipt can say:

> *"This unit of work was performed on hardware profile H under conditions C, with proof P and VerifyPrice statistics V."*

Monetarily, that matters because:

- Profiling makes **hardware risk priced** instead of hidden. Credits from a weak profile H' can trade at a discount.
- It lets different actors pick their risk tolerance: some will only hold Work Credits linked to fully open cores; others will accept mixed profiles in exchange for lower cost or higher performance.

Layer 0's job is not to tell everyone what risk to take; it is to make the **risk legible and instrumentable**.

17.4 The Layer 0 Feasibility Ladder

A common objection to Layer 0 is: "Open silicon and sampled supply chains sound like moonshots. What is actually achievable *this decade* given real-world fabs, opaque GPU stacks, and geopolitical constraints?"

The answer is a **graded trust ladder**. Layer 0 does not require perfection on day one; it requires **measurable progress and falsifiable claims** at each grade. Higher grades provide stronger guarantees; lower grades are acceptable for less sensitive workloads, with telemetry that detects when you're relying on weaker grades.

Grade	Name	What It Means	Timeline
L0-A	Best Available Today	Attestation + multi-party audits + reproducible benchmarking + diversity	Now
L0-B	Sampled & Bounded	Lot sampling + imaging + side-channel budgets + independent lab inspections	1–3 years
L0-C	Partial Open	Open RTL for critical components; open host/router firmware; proprietary accelerators admitted only for workloads with cheap algebraic verification (see below)	2–5 years
L0-D	Fully Open	Fully open designs + open PDKs + multi-fab production	Now, for RoT/RNG/metering/network-controller components at 130-

Table 17.1: Layer 0 Feasibility Ladder

Scope of the top two grades, stated plainly. Two cells in that table are easily read as engineering timelines, and they are not. First, L0-D. Open PDKs exist today only at mature nodes—SkyWater SKY130, GlobalFoundries GF180MCU, IHP SG13G2, i.e. 130–180 nm—and nothing resembling an open PDK exists, or is on any announced path, at the leading-edge nodes accelerator-class silicon requires. L0-D is therefore *viable now* for the components whose job is trust rather than throughput: roots of trust, RNGs, power and capacity meters, network controllers, and the attestation logic that binds them. For provers and verified-compute accelerators, L0-D is a research goal with no known path, and this document does not assign

it a date. Second, L0-C's "accelerators wrapped with proofs." Wrapping a proprietary accelerator so that its internals need not be trusted works only where the output can be checked far more cheaply than it was produced: Freivalds-class randomized checks on matrix products [Freivalds 1977], Merkle-committed outputs against sampled re-execution, and workloads with a native algebraic verifier. Generic verifiable GPU compute—proving that an arbitrary kernel ran correctly at accelerator speed—is an open research problem, and L0-C admits proprietary accelerators only for the workload classes where the cheap check exists. Where the thesis elsewhere speaks of "proof-wrapped accelerators," it is to be read with this restriction.

Key principles:

1. **Grades are explicit.** Every hardware profile is tagged with its Layer 0 grade. Work Credits, VerifyPrice dashboards, and SLA tiers reference these grades so users know what trust assumptions they are accepting.
2. **Telemetry detects reliance on weaker grades.** If 80% of proving capacity is L0-A (vendor-attested) and only 5% is L0-C (partial open), that concentration is visible in dashboards. Users can decide whether to hold Work Credits tied to such a distribution.
3. **Higher grades earn lower risk premiums.** Markets should price Work Credits from L0-D profiles more favorably than L0-A profiles, creating economic gravity toward openness as it becomes available.
4. **Migration paths are first-class.** When L0-B or L0-C options become viable, there are documented procedures to migrate workloads off L0-A profiles without catastrophic downtime.
5. **Failure modes are bounded.** If a specific grade is compromised (e.g., a TEE used in L0-A profiles is broken), the impact is contained to that grade. Higher-grade capacity continues to function; affected Work Credits are risk-flagged, receive collateral haircuts, or face prospective ineligibility.

17.4.1 Policy Hooks: Economic Treatment by Grade

The following table specifies how each L0 grade affects Work Credit eligibility, pricing, and collateral treatment. These are reference parameters; implementations may adjust within bounds.

Interpretation:

- **Issuance weight:** No grade earns *more* than one Work Credit per standardized unit of committed capacity. A Work Credit is a typed claim on reserved service (Chapter 12); issuing 1.1 or 1.2 claims per unit committed would make the claims exceed the deliverable work and turn the instrument into a subsidy token, in conflict with the issuance bound of Section 17.9. Lower grades are haircut *below* baseline (fewer claims than capacity committed, so the outstanding stock stays over-backed); the incentive for open-hardware investment is carried by collateral grade, routing priority, and the VerifyPrice premium tier, none of which inflate the claim stock.

Grade	WC Eligibility	Issuance Weight	Collateral Haircut	VerifyPrice Tier
L0-A	All workloads; capped at 60% of total issuance	0.9× (below baseline)	0% (if diversified) to 15% (if concentrated)	Reported separately; flagged if >70% of measurements
L0-B	All workloads; no issuance cap	1.0× (baseline)	0%	Standard reporting
L0-C	All workloads; eligible for "open-profile" premium tier	1.0× (baseline; premium carried by collateral grade and routing priority, not by issuance)	0%; may qualify as "pristine collateral"	Premium tier; highlighted in dashboards
L0-D	All workloads; highest openness grade	1.0× (baseline; premium carried by collateral grade and routing priority)	0%; pristine collateral (with Gold assurance, Section 17.5.4)	Open-reference tier; used as reference baseline

Table 17.2: Economic treatment by L0 grade.

- **Collateral haircut:** When WC are used as collateral (e.g., for staking, LP positions, or DeFi), lower grades may face haircuts. L0-A collateral is discounted if a single vendor/TEE dominates.
- **VerifyPrice tier:** Dashboards stratify measurements by L0 grade. If most verification runs on L0-A hardware, this is flagged as concentration risk. The ladder grades *openness*; the **Gold assurance tier** of Section 17.5.4 grades *sampling evidence* and is orthogonal to it: a profile can be L0-B and Gold, or L0-D and not yet Gold.
- **Issuance cap (L0-A only):** To prevent over-reliance on vendor-attested hardware, L0-A profiles are capped at 60% of new issuance. Excess work at L0-A earns priority for future slots but no additional Work Credits. The incentive consequence is stated rather than left to be discovered: at the cap, an L0-A operator's marginal delivered unit is unpaid in Work Credits, so operators at the margin will decline L0-A work or migrate hardware grade — which is the cap's purpose, and the reason the cap is published in advance rather than applied ex post.

17.4.2 Collateral Haircut by Modality and Grade

The collateral haircut in the policy-hooks table above is indexed on openness grade alone. Section 22.2.1 shows that the more fundamental axis is the verification modality of the position being pledged: an M3-verified position carries irreducible delegated-trust exposure whatever its hardware grade, and an M1-verified position on modest hardware may carry less. The two axes are therefore indexed jointly, and the joint table is the one collateral integrations read; the grade-only column above is its M1 row.

Openness grade	**M1 succinct** (Tier A)	**M2 probabilistic** (Tier B)	**M3 replicated + audit** (Tier B)	**M3 attestation-only** (Tier C)
L0-C / L0-D with Gold assurance	0% — pristine	10%	25%	Ineligible
L0-C / L0-D, not yet Gold	0% — standard	10%	25%	Ineligible
L0-B	0% — standard	10%	25%	Ineligible
L0-A, diversified	0–5%	15%	30%	Ineligible
L0-A, single vendor $> 40\%$	15%	25%	40%	Ineligible

Table 17.3: Reference collateral haircuts on pledged Work Credit positions, by openness grade and verification modality. Haircuts discount what is *pledged*; they are distinct from the issuance weights of Section 25.6.5, which discount what is *minted*. The modality increment (read across a row) exceeds the grade increment (read down a column) at every cell except the concentrated L0-A row, which is the claim of Section 22.2.1 in numbers: a committee's word costs a lender more than a closed die does. Tier C is collateral-ineligible at every grade (Section 22.10). Governance parameters, adjustable under the same disclosure rules as the thresholds above.

Metric	**Threshold**	**Consequence if Breached**
L0-A share of issuance	>60%	New L0-A issuance paused until ratio falls
Single vendor share within L0-A	>40%	Affected profiles moved to "watch list"; 5% collateral haircut
Sampling coverage for L0-B	<80% of production lots sampled per quarter (at tier minimum unit rates, ≥5 units per sampled lot; Section 17.5.4)	Grade downgraded to L0-A until coverage restored
Open-source audit currency (L0-C/D)	>12 months since last audit	Grade downgraded one level

Table 17.4: Quantitative thresholds for L0 grades.

17.4.3 Quantitative Thresholds (Reference Design)

Read the first row plainly, because the number is doing more work than it appears to. Three-fifths of the world's proving running on vendor-attested chips is the state the cap forbids: every *attestation* in that world—every claim about which machine did the work, on what hardware, at what energy cost—is a signature away from being repudiated by whoever holds the signing key (proof soundness itself does not depend on the prover's hardware, per the sidebar in Section 17.1, which is exactly why the exposure is in the attestations and not the proofs), and the stack's central promise—that anyone can verify without trusting a platform—quietly reduces to trusting three vendors in sequence. The thresholds exist to make that drift expensive before it is complete. Each is a tripwire that fires while the concentration is still a governance problem, rather than after it has become a settlement failure. Breaching one does not crash anything; it pauses issuance, raises a haircut, downgrades a grade. The damage is contained by making the drift itself the first thing to pay.

These thresholds are governance parameters, adjustable via protocol upgrade with supermajority. Changes are announced 90 days in advance and visible in dashboards.

17.4.4 Pragmatic Starting Point (L0-A)

Today's stack can achieve L0-A by combining:

- **Multi-vendor diversity:** No single chip family or TEE dominates more than 40% of critical capacity, aligned with the single-vendor share trigger of Section 17.4.
- **Third-party attestation audits:** Independent labs verify that attestation claims match device behavior.
- **Reproducible firmware and benchmarks:** Host, router, and metering firmware is built reproducibly from published source. Accelerator firmware is not: GPU and ASIC firmware is vendor-signed and closed, and cannot be rebuilt by anyone outside the vendor. For accelerators the requirement is therefore *hash-pinned and version-attested*: the profile names the exact firmware digests it accepts, devices attest the running version, and the residual—that the vendor's binary is trusted as shipped—is recorded in the profile dossier as a named trust assumption rather than left implicit. Benchmark results are publicly verifiable.
- **Side-channel measurement:** Known side-channel leakage is measured and published as a "leakage budget" per profile.
- **Cryptographic agility:** Profiles document which primitives can be upgraded via firmware vs. require hardware swap.

This is not trustless. It is *trust-bounded and measured*. The residual trust is explicit: "We trust these vendors, these auditors, and this sampling methodology, and here is the evidence."

17.4.5 The Path Forward

- **L0-A → L0-B:** Fund independent lot-sampling programs. Publish inspection results. Build statistical models of detection confidence.
- **L0-B → L0-C:** Invest in open RISC-V cores, open RNG designs, open host and router firmware stacks. Wrap proprietary accelerators with proof interfaces *for the workload classes where a cheap check exists*—Freivalds-class MatMul verification [Freivalds 1977], Merkle-committed outputs with sampled re-execution—so that the accelerator's internal logic need not be trusted for those workloads. For workloads without such a check, the accelerator remains a trusted component and the profile says so.
- **L0-C → L0-D:** Move the trust-critical, low-throughput components—roots of trust, RNGs, meters, network controllers—onto open PDKs at the mature nodes where those PDKs exist today, manufactured at multiple fabs in multiple jurisdictions. For accelerator-class silicon this step has no known path; funding open-PDK research at advanced nodes is the honest description of the work, and it is research, not a roadmap item with a date.

17.4.6 Why This Prevents "Layer 0 Is Impossible, Therefore Thesis Fails"

The thesis does not require L0-D today. It requires:

1. Clear grading of what trust assumptions each profile carries.
2. Telemetry that makes those assumptions visible.
3. Economic and technical paths toward stronger grades over time.
4. Failure modes that degrade visibly, not silently.

If these four conditions hold, Layer 0 becomes a *progress metric* rather than a moonshot prerequisite. The stack can operate today at L0-A while building toward L0-C/D, and users can decide what risk they are willing to accept at each stage.

17.5 Concrete Components of Layer 0

Layer 0 is not a single device. It is a **bundle of practices and mechanisms** that give higher layers a surface to stand on.

17.5.1 Open Designs Where Possible

The gold standard is **open cores and toolchains**: open RTL or microarchitectures for CPUs/accelerators, open PDKs where geopolitical conditions allow, and reproducible build infrastructure.

Where open options exist, they are **first-class citizens** in hardware profiles. On the margin, this creates **economic gravity**: as open hardware matures, Work Credits tied to open profiles

should command a lower risk premium.

17.5.2 Attested Randomness and Entropy

Randomness is the hidden spine of consensus and proofs. A biased RNG can leak signing keys, predict leader election, and make PoUW "fairness" an illusion.

Layer 0 requires:

- Hardware RNG designs that are **documented and testable**.
- **Attested entropy tests**: periodic statistical test suites.
- **Blended randomness**: mixing hardware entropy with VRFs, commit-reveal, and cross-machine aggregation.

17.5.3 Attestation Without Priesthood

Modern hardware ecosystems push TEEs and attestation as the answer to everything. Used naively, they simply move "trust the vendor" into "trust the vendor's signing key."

Layer 0's posture:

- TEEs and hardware attestation are **useful tools**, not **root of trust**.
- Attestations should be **wrapped and sampled**, not taken as gospel.
- Devices emit **local attestations** wrapped in **SNARKs or STARKs** where possible.

The title of this subsection claims less than it may appear to, and the limit should be stated where the mechanism is introduced rather than discovered later. A succinct proof over a vendor attestation proves exactly one thing: that a valid signature under vendor key K exists over the attested measurement. It does not prove that the measurement is true, that the device is what K says it is, or that K has not been misused. The vendor key remains the root of trust for the attestation's content; wrapping does not move it. What wrapping delivers is **compression** (one proof over many attestations, verifiable on any chain at fixed cost) and **selective disclosure** (revealing that a device belongs to an accepted profile without revealing which device). Both are worth having. Neither is trust reduction. Trust reduction in this stack comes from two places only: the lot-sampling program of Section 17.5.4, which checks whether devices behave as attested, and open design (L0-C/D), which lets the attestation logic itself be inspected. "Without priesthood" is therefore earned by sampling and openness, with wrapping as the transport that makes their results cheap to consume; it is not earned by the wrapping alone, and the same limit applies wherever this document speaks of ZK-wrapped quotes (Section 20.5, Section 20.8).

17.5.4 Lot Sampling and Destructive Audits

Because no one can open every chip, Layer 0 leans on **lot sampling**:

- For each hardware profile, a fraction of units are randomly selected for **deep inspection**:

decapping, imaging, side-channel probing.
- Results are published as part of the profile's dossier.
- A protocol-level **Layer 0 Assurance Fund**, the largest line of the assurance budget in the fee routing (10% of gross fees, Section 25.6), finances sampling.

This is how Layer 0 turns "we hope the vendor is honest" into:

> *"We have inspected a statistically relevant sample of this profile, and here are the findings and residual risks."*

Sampling Economics and Sufficiency

Critics will ask: "How much sampling is enough? Who pays? What confidence do we actually get?" The honest answer begins by conceding the premise: sampling does not produce certainty, and no number in this subsection should be read as though it does. What sampling produces is a *published, priced, third-party* bound on how much of a fleet is compromised *by the defects the tests can see*—an instrument with two distinct commercial pedigrees, and it is worth keeping them apart. The sampling precedent is acceptance sampling: MIL-STD-105 and its civilian successor ANSI/ASQ Z1.4 [DoD 1989] have governed lot-by-lot attribute sampling for industrial procurement since the 1950s, and pharmaceutical lot-release testing under FDA and ICH guidance accepts or rejects a production lot on the basis of a drawn sample rather than a full inspection. Both accept a stated consumer's risk in exchange for an affordable inspection; both publish the sampling plan so the risk can be computed rather than asserted. The publish-and-underwrite precedent is Lloyd's Register, and it is a different thing: Lloyd's classification was a per-vessel survey, not a sample, and it is cited here not for its statistics but for its market structure—a third-party grade, paid for by the industry that needed it, published, and independent of the shipowners, which underwriters then priced off. Marine insurance priced off Lloyd's grade the way collateral markets here price off L0 grades. The statistics below come from the acceptance-sampling tradition; the institutional form comes from Lloyd's; neither is claimed to be the other.

Sample rate model (reference design):

Units and lots. The table's minimums are counted in *units*; the coverage thresholds of Section 17.4 and the Gold assurance tier below are counted in *lots*. The two are reconciled by a per-lot floor: every lot selected for destructive inspection contributes at least **5 units**, drawn at random from within the lot. A profile's quarterly sampling obligation is therefore the larger of two numbers—the tier's unit minimum, and five times the number of lots its coverage threshold requires. A Tier 1 profile shipping twelve lots a quarter at 80% coverage owes $\max(50,\ 5 \times 10) = 50$ units; the same profile shipping forty lots owes $\max(50,\ 5 \times 32) = 160$.

Profile Criticality	Definition	Min Sample Rate	Funding Source
Tier 1 (Critical)	Profile in the top 5 by WC issuance OR above 10% of total capacity; a profile meeting either condition is Tier 1	≥50 units/quarter	Protocol assurance budget (the 10% assurance share of Section 25.6)
Tier 2 (Standard)	Profile with 1–10% of capacity	≥20 units/quarter	Protocol assurance budget
Tier 3 (Emerging)	New profiles in probation; <1% of capacity	≥5 units/quarter	Profile sponsor (operator or vendor)

Table 17.5: Sample rate model by profile criticality.

The unit minimum is a floor for small-volume profiles, not a ceiling for large ones.

Statistical sufficiency: For a profile with N deployed units, sampling n units provides confidence that:

$$P(\text{detect} \mid \geq k\% \text{ compromised}) \;\geq\; 1 - \left(1 - \sigma\,\frac{k}{100}\right)^{n}, \qquad 0 < \sigma \leq 1,$$

where σ is the **test sensitivity**: the probability that the inspection regime actually applied to a sampled unit flags that unit, given that it is compromised. Setting $\sigma = 1$ recovers the textbook bound, and the textbook bound is the wrong number to publish, because σ is not one and for the most dangerous defect class it is not known.

Two assumptions carry the bound. The first is **uniform prevalence**: compromise is modeled as spread evenly across the deployed population, so that a random sample of n units meets compromised units at the background rate. Concentrated compromise—an entire production lot backdoored while others are clean, or 20% of lots attestation-only and never destructively inspected—defeats it, because prevalence in the sampled lots is then not prevalence in the fleet. This is why lot selection is randomized and stratified rather than convenience-sampled, and why the attestation-only residual is a published quantity rather than a silent one. The second is **sensitivity**, and it is where honesty costs the most. For counterfeit dies, remarked or binned-down parts, substituted firmware, and metering fraud, σ is high: these defects are visible to electrical test, decapping and optical inspection, firmware-hash comparison, and cross-metering, and a competent lab finds them when it looks. For stealthy hardware trojans in leading-node accelerators—dopant-level modifications of the kind Becker et al. demonstrated at CHES 2013 [Becker et al. 2013], which alter transistor behavior without changing any metal layer an imaging pass would see—σ against a quarterly sampling program is unknown and probably low. Full-die reverse engineering at 5–7 nm is not something an assurance fund buys every quarter, and a trojan designed to evade it is designed by someone who knows that.

The sampling program's primary, defensible targets are therefore the high-σ classes: counterfeit dies, binned-down parts, firmware substitution, and metering fraud. Against those it produces a bound worth publishing. Against dopant-level trojans in advanced-node silicon it produces a deterrent and a paper trail, not a bound, and the profile dossier says which of the two it is offering.

For Tier 1 profiles ($n = 50$, $k = 5\%$), the detection probability is ∼92% at $\sigma = 1$ and ∼72% at $\sigma = 0.5$; at the low, unknown σ appropriate to stealthy trojans the bound is not quoted because it would be a number without a basis. Published bounds carry the σ assumed and the defect class it applies to. This is not certainty—it is **bounded uncertainty**, documented and priced, and bounded only for the defect classes the tests can see.

The converse matters just as much and must be stated with equal care. With n samples and

zero findings, the 95% upper confidence bound on the prevalence of *detectable* compromise is approximately $3/(\sigma n)$ (the rule of three, sensitivity-adjusted). Fifty clean samples at $\sigma = 1$ bound undetected compromise at ∼6% (the exact bound $1 - 0.05^{1/50}$ is 5.8%), not at zero; at $\sigma = 0.5$ the same fifty clean samples bound it at ∼12%. Lot sampling supports *graded confidence*, never exoneration, and no claim in this thesis should be read as the latter.

The Gold Assurance Tier

It is tempting to identify "Gold tier" with grade L0-D. That is a category error: the feasibility ladder grades *openness*—how much of the design can be inspected—while the conditions below are all about *sampling evidence*, and say nothing about RTL or PDKs. The two are kept orthogonal. **Gold assurance** is a tier a profile earns through its inspection record; it is attainable at openness grade L0-B or above, and L0-D remains the openness grade, earned by design and PDK disclosure alone. A profile can be L0-B and Gold (a closed design with an exemplary, heavily sampled supply chain), or L0-D and not yet Gold (a fully open design whose fleet has not accumulated four clean quarters).

A profile achieves Gold assurance when:

1. ≥95% of production *lots* are sampled each quarter at or above the tier minimum unit rate and the per-lot floor of ≥5 units, with the remaining lots covered by supply-chain attestation. This is deliberately stricter than the ≥80% floor that keeps a profile at L0-B (Section 17.4); a threshold that merely keeps a grade cannot also be the criterion for the highest assurance tier.
2. Zero critical findings (backdoors, RNG bias >threshold, firmware or die substitution, metering discrepancies) in the last 4 consecutive quarters.
3. Side-channel leakage remains within published budget under independent testing.
4. Inspections in each of those four quarters were conducted by at least 2 independent labs, and no single lab performed more than 70% of the units sampled.

Coverage is counted in *lots*, not deployed units. Inspections are destructive and expensive; the lot is the auditable unit a sampling program can afford to buy. "95% of deployed units" would describe a near-total destructive audit of the fleet, which no assurance budget can fund and no fleet can survive; the reference design does not propose it and this document should not be read as though it does.

Gold assurance unlocks:

- For L0-C and L0-D profiles, the "pristine collateral" designation in DeFi integrations. Pristine requires both: inspectable logic (openness ≥ L0-C) *and* Gold assurance; the grade table's "may qualify" for L0-C is this condition.
- Release from probation weights (Section 17.11) and eligibility as a reference baseline in VerifyPrice stratification.
- For L0-B profiles, an "assured" flag in dashboards and profile dossiers. L0-B with Gold

assurance does *not* earn pristine status; a closed design remains closed however thoroughly it is sampled.

Issuance eligibility and the 0% collateral haircut are consequences of the openness grade (L0-B and above), not of Gold assurance; the grade table governs them, and Gold adds nothing to them. "Pristine collateral" as a designation therefore means openness grade L0-C or L0-D *with* Gold assurance throughout this document.

Funding mechanism: The fee routing of Section 25.6 allocates 10% of gross fees to an assurance budget; the **Layer 0 Assurance Fund** is that budget's largest line, alongside security review and audit. The fund:

- Contracts with independent hardware security labs.
- Publishes RFPs for sampling campaigns.
- Maintains a public registry of inspection results.
- Is governed by a multisig of hardware security researchers (not protocol developers).

Transparency: All sampling results are published within 30 days. Raw data (images, test logs) is archived and available for independent verification. If a lab's findings are disputed, a second lab can be commissioned for arbitration.

17.5.5 How This Opens New Economies

Verifiable machines do not merely make today's cloud slightly less opaque. They unlock whole categories of economic arrangement that are not viable when hardware is a black box.

Verifiable cloud and compute co-ops. If you can spin up a rack of open-design TEEs or accelerators and have them produce attestations and ZK receipts that anyone can verify, then:

- A small data center in Nairobi or Reykjavík can sell the same class of "trusted inference" or "trusted proving" as a hyperscaler in Virginia.
- A co-op of households can pool "AI appliances" in their basements and earn by running verified work for others.
- Regulators and enterprises can enforce compliance through proofs and telemetry, not vendor logos.

Compute ceases to be a winner-takes-all brand game and becomes a commodity with open admission.

Civic infrastructure that does not depend on one vendor's conscience. Voting machines, digital ID kiosks, public-health dashboards: today they are RFPs to a short list of contractors. With open hardware and proof-wrapped attestation, you can build ballot boxes and ID hardware whose entire stack is open to public inspection and require that each device emit public

proofs of correct behavior. Democratic legitimacy becomes a property of math and sampling, not of which vendor's logo sits on the plastic.

Tamper-evident telemetry. As the physical world fills with sensors and actuators, the ability to sell tamper-evident telemetry becomes critical. An environmental sensor built on an open profile, with attested firmware and ZK-wrapped readings, can sell CO_2 or temperature data as evidence into climate, insurance, and industrial hedging markets—not just as numbers on a dashboard.

Hardware-native credit instruments. Once machines themselves are verifiable actors, capacity becomes something a credit instrument can be written against: "miner-notes" backed by the future output of a specific open-design proving farm, or project paper whose coupon is denominated in SLA-backed capacity—X proofs per second, Y verified FLOPs, Z private swaps per month—all attested by open hardware meters. These are instruments of the labeled, non-monetary **Open Duration Warehouse** (Section 34.5.1): they are project credit, carrying construction risk, counterparty risk, and time risk, underwritten by loss-bearing holders who can sit through marks, and their novelty is only that the covenant is checkable by receipt rather than by site visit. They are explicitly *not* the monetary object. The base asset carries no coupon, capacity-denominated or otherwise; the moment it did, proofs would have become the bond and duration-neutrality—which this thesis treats as a repression-resistance requirement, not a preference (Section 3.3)—would have been undone through the protocol door. Verifiable machines make capacity credit underwritable. They do not, and must not, make it money.

Open profiles manufactured at multiple fabs in multiple jurisdictions turn chip supply into a multi-polar fabric instead of a single chokepoint. Neutral money needs neutral hardware; verifiable machines produced on a diversified manufacturing base make that a reachable design goal rather than a slogan.

17.6 Verifiable Power: Energy as First-Class Input

If Work Credits are energy-anchored claims on triad work, then **power** is not just an environmental footnote; it is an input to service delivery. It does not back the base asset or make the credit monetary.

Layer 0 treats power in three ways:

1. **Measurement.** Prover farms track power draw, mapping between workloads and power, and local generation vs. grid intake. These feed into **Work Credit metadata** and **VerifyPrice** models.
2. **Resilience.** Micro-grids or backup generation for critical infrastructure; geographic and jurisdictional dispersion of power sources.

3. **Policy hedging.** Transparent power telemetry allows answering "how much of this is actually training/inference/proof that humans pay for?"

For Work Credits, this means credits can carry **energy provenance tags** and markets can price credits differently based on energy profile.

17.7 Facility Capacity Receipts: Beyond Energy

Facility Energy Receipts (FERs, Appendix A.2) make energy legible, but verified digital infrastructure depends on more than kWh. A facility can have cheap energy and still be fragile if it sits behind a congested interconnect, depends on a water-constrained cooling system, runs in a jurisdiction with curtailment risk, or relies on a single transformer, fuel source, cloud vendor, or hardware profile.

Layer 0 therefore extends FERs into **Facility Capacity Receipts**. An FCR records the physical capacity envelope behind a unit of verified work in ten fields (Table 17.6): grid node, interconnection class, outage history, backup duration, fuel mix, cooling dependency, transformer redundancy, hardware profile mix, jurisdictional risk score, and a confidence grade on the receipt itself. Curtailment exposure is carried by the grid-node and jurisdictional-risk fields; water intensity by the cooling-dependency field. The list is ten fields, not more, and the same ten are what Section 17.8 aggregates.

Facility Capacity Receipt (FCR)

A signed, auditable receipt that extends Facility Energy Receipt data with infrastructure resilience, grid, cooling, jurisdictional, and hardware-diversity metrics.

Table 17.6: Suggested FCR fields

Field	Why it matters
Grid node / balancing authority	Captures congestion and curtailment risk
Interconnection class	Captures firmness of power access
Outage history	Captures reliability
Backup duration	Captures blackout survivability
Fuel mix	Captures geopolitical/commodity exposure
Cooling dependency	Captures water and heat constraints
Transformer redundancy	Captures single-point grid failure
Hardware profile mix	Captures compute centralization
Jurisdictional risk score	Captures seizure, power rationing, sanctions risk
FCR confidence grade	Allows risk-weighted Work Credit issuance

A Work Credit backed by resilient, diversified, verifiable capacity is economically different from one backed by cheap but brittle compute. The market should be able to price that differ-

ence. The stack should not hide physical fragility behind cryptographic elegance. If FCR data is unavailable or unverifiable for a facility, credits from that facility should receive issuance caps or risk haircuts (Section 25.6.5); Section 22.8 formalizes the cost of auditing FCR claims as **Physical VerifyPrice**.

17.8 Sovereign Optionality: Aggregating FCR Into One Exposure

Table 17.6 gives an allocator ten fields per facility. That is the right raw material and the wrong interface. Nobody underwrites a monetary asset by reading ten columns across a thousand facilities, and a field-by-field presentation makes it easy to be locally reassuring while globally fragile: every facility can look acceptable while the network sits inside two balancing authorities and one jurisdiction.

A single aggregate exposure is therefore defined over fields the stack *already collects*. This is deliberately not new telemetry. It imposes no additional reporting burden; it changes how existing FCR data is read.

Sovereign Optionality ($\mathcal{O}_s$)

A capacity-weighted composite of substitutability attributes—fuel diversity, firmness, interconnection, jurisdictional dispersion, and hardware mix—computed from Facility Capacity Receipt fields. It *correlates with* the number of independent physical pathways by which verified work can continue under disruption, but does not literally count them; the pathway-count question is answered by Delivered Verified Capacity (Section 17.9).

17.8.1 Adaptation of Scale

The concept is borrowed from state-level energy strategy [Doomberg 2026a], where analysts assess how many independent ways a sovereign can obtain electricity, fuel, and industrial feedstock when ordinary market relationships fail. The transposition to a proving network is close but not exact: the state-level analysis includes industrial and logistical throughput, and proving has a conversion analogue of its own.

A proving facility has a measurable conversion chain:

energy/fuel → firm electricity → transformer/switchgear → cooling/water → hardware → network → wor

It is not manufacturing throughput, but it is conversion throughput. Any edge can become the binding cut [Doomberg 2026b], and a weighted resilience score can conceal a zero-flow series bottleneck. Sovereign Optionality therefore remains the substitutability measure; Delivered Verified Capacity below supplies the missing throughput measure.

17.8.2 Mapping to Existing FCR Fields

Table 17.7: Sovereign optionality components and their FCR sources

Term	Component	FCR field(s) used
D	On-site availability	Fuel mix; local generation vs. grid intake (Section 17.6)
F	Conversion flexibility	Fuel mix diversity: Shannon entropy over independent fuel-source shares, normalized by $\log(n_{\max})$ so $F \in [0,1]$, matching the other components' $[0,1]$ normalization. $n_{\max}$ is the number of fuel-source classes in the published FCR fuel-mix taxonomy (a fixed, versioned constant shipped with the telemetry spec), not the number of sources a given facility happens to use; otherwise a single-source facility would score $F = 1$ trivially
I	Reserve depth	Backup duration
G	Grid firmness	Interconnection class; outage history; transformer redundancy; grid node
V	Supplier & technology diversity	Hardware profile mix
C	Coercion exposure *(subtracted)*	Jurisdictional risk score; curtailment exposure

Every input already exists. Note in particular that the jurisdictional risk score of Table 17.6 is already specified to capture "seizure, power rationing, sanctions risk." The stack has been collecting the coercion term all along without aggregating or acting on it.

17.8.3 Facility and Network Form

For a facility f, with each component normalised to $[0,1]$:

$$\mathcal{O}_s(f) \;=\; w_D D_f + w_F F_f + w_I I_f + w_G G_f + w_V V_f \;-\; w_C C_f, \qquad w_i \geq 0, \quad \sum_{i \neq C} w_i = 1,$$

so that the positive part of the index lies in $[0,1]$ and $\mathcal{O}_s(f) \in [-w_C,\, 1]$: the declared range is fixed by the constraint, and a published w_C tells the reader exactly how far below zero a maximally coerced facility can score. Without the constraint the weights could be rescaled to move every facility's score without changing any input, which is the failure mode Section 17.8.4 exists to prevent.

The network-level figure is the capacity-weighted mean across facilities, less a concentration penalty. Let κ_f be facility f's share of verified capacity. H is the **average of two capacity-share Herfindahl–Hirschman indices**, one computed over the jurisdiction partition of the fleet

and one over the balancing-authority partition. That single convention is pinned here; no other aggregation of the two partitions is used anywhere in this document:

$$\mathcal{O}_s^{\text{net}} \;=\; \sum_f \kappa_f \, \mathcal{O}_s(f) \;-\; w_H H, \qquad H \;=\; \tfrac{1}{2}\big(H_{\text{jurisdictions}} + H_{\text{balancing authorities}}\big)$$

The penalty term is what prevents the local-reassurance failure. A network of individually excellent facilities that all sit behind the same interconnect or inside the same legal regime has one pathway, not many, and the weighted mean alone would not say so.

One overlap in this form is deliberate and should be visible rather than silent: jurisdictional exposure enters twice, once per-facility through $w_C C_f$ (siting-level risk: seizure, rationing, sanctions) and once network-level through $w_H H$ (concentration of the fleet inside few jurisdictions or balancing authorities). Because the two terms price different failure modes at different scopes, both are retained; the weights w_C and w_H are therefore *published jointly and reviewed annually*, so the overlap is a documented modelling choice an allocator can see and reweight, not a hidden double count.

Relationship to existing metrics. H overlaps deliberately with the dispersion metrics already used for verification monoculture and enclosure risk (Section 30.6, Red Lines 3 and 11). $\mathcal{O}_s^{\text{net}}$ should consume those existing measurements rather than introduce a parallel concentration statistic. Where they disagree, that disagreement is itself a finding.

17.8.4 Governance of the Weights

The weights w_i are judgement, not measurement, and an index whose weights can be quietly retuned is a marketing instrument. Three constraints follow:

- **Published and versioned.** Weights ship with the telemetry spec; changes are proposals with rationale, not silent parameter updates.
- **Independently reproducible.** A third party holding the same FCR data must be able to recompute $\mathcal{O}_s$ and obtain the same number.
- **Subject to capture review.** Weight-setting falls under the same telemetry-capture constraints as the rest of the scoreboard (Section 30.6, Red Line 4). The party being scored must not set the scoring function.

Report $\mathcal{O}_s$ alongside its component vector, never as a bare scalar. A single number is what an allocator wants and also what conceals which substitutability attribute is missing; publishing both is the only honest form.

17.8.5 What $\mathcal{O}_s$ Is For

The index does real work in two places, and is otherwise decoration:

1. **Risk haircuts.** It extends issuance discipline from "we cannot see this facility" to "we can see it clearly and it is fragile" (Section 25.6.5).
2. **Falsification.** It supplies the measurable condition for the energy-sovereignty red line (Section 30.6).

Absent those two hooks, $\mathcal{O}_s$ would be another dashboard number. With them, it is a constraint that can bind against issuance and a threshold that can retire the thesis.

17.9 Delivered Verified Capacity: Surviving Service Flow

Picture a city's water system. The reservoir level is the number a brochure quotes; what comes out of the tap is what the city actually drinks. Between the two runs a network of mains, valves, and pumping stations, and any one of them—not the reservoir—can be the thing that decides whether water arrives. DVC is the tap, not the reservoir. Nameplate capacity says how much could flow if nothing failed; Delivered Verified Capacity says how much does flow once a scenario has removed, throttled, or delayed the edges it removes, throttles, or delays. The narrowest main in the network — the minimum cut, in the single-workload case — is the only number that matters while it binds. This is why the measure that follows is a flow computation rather than a capacity sum: the sum flatters, the cut confines. And it is why substitution latency enters the arithmetic at all—a backup reservoir that takes three weeks to route is not redundancy for a fire that started today.

Delivered Verified Capacity (DVC)

For canonical workload W, interval t, and disruption scenario s, Delivered Verified Capacity is the maximum usable service *flow per unit interval* through the complete energy-to-settlement network after scenario-specific edge capacities, substitution latency, and common-cause failures are applied. It is a rate, not a stock.

DVC is defined here but deliberately parameterized by objects defined later: the workload W is drawn from the canonical workload registry of Layer 4, and the sink is usable *settled* service in the sense of Layer 5 settlement. The same W and the same sink are used throughout this Part and the two economic layers; nothing downstream redefines them.

Let $G_{W,t,s} = (V, E)$ be a directed capacity graph whose source nodes are available energy and fuel and whose sink is usable settled service. Each edge e carries surviving capacity $c_e(W, t, s)$ in a common workload-specific service unit. Then

$$\mathrm{DVC}(W, t, s) = \mathrm{MaxFlow}(G_{W,t,s}, c(W, t, s)) \,,$$

and the constraints that bind at the optimum identify the bottleneck that actually limits delivery. The unit is not a generic FLOP: it is a settled, independently verifiable unit of the canonical

workload at the promised tier. Because DVC is a flow, the quantity comparable to it is issuance over an interval, and the bound below carries the interval explicitly.

Units change along the chain; the flow must be told so. A plain max-flow assumes every edge is measured in the same unit, and the chain of Section 17.8.1 is not: fuel enters in MWh-equivalent, electricity leaves the switchgear in kW, the hardware node emits proofs per second, and the sink counts settled units of W. The correct object is therefore a **generalized flow** in which each node carries a gain factor—the kW$\to$proofs/s conversion is a property of the hardware class sitting at that node, published in its profile as a measured efficiency at the promised tier—and the flow is conserved only after the gain is applied. Equivalently, the graph must be built so that every energy edge terminates at exactly one hardware class, with the conversion pinned at that node, and no edge is permitted to carry "capacity" in a unit its successor cannot consume. Either construction recovers a well-defined maximum flow; a graph that mixes kW and proofs/s on edges feeding a common node does not, and a DVC quoted from such a graph is not a number. A second consequence follows for issuance across workloads. The energy, cooling, transformer, and network edges of a facility are shared among every W it serves, so computing $\mathrm{DVC}(W, t, s)$ independently for each W and summing counts the same transformer once per workload. Issuance across several workloads therefore requires either a **multi-commodity formulation**—all workloads routed simultaneously through the shared edges, each edge's capacity binding on their sum—or a **published shared-capacity allocation rule** that assigns each shared edge's surviving capacity to workloads in declared fractions before any per-workload flow is computed. Which of the two a constitution adopts is a governance parameter; that it must adopt one of them is not. The textbook min-cut reading does not survive either construction intact. Maximum flow equals minimum cut for a single commodity on a lossless graph; with gain factors, or with several workloads sharing edges, that equality fails in general, and a cut is no longer a certificate of the bottleneck. What survives is the **binding constraint**: the solver reports which capacities are saturated at the optimum and how much delivery each would add if relaxed, and that is the bottleneck the boards publish. Ordinary max-flow/min-cut remains the reference model for one workload on a graph whose units are pinned at the nodes.

Scenario specificity. A benign-state graph and a drought, fuel-shock, transformer-loss, chip-embargo, network-partition, stablecoin-depeg, or coordinated-policy graph are different objects. Scenario s changes edge capacities and may remove nodes or entire dependency classes. Publishing only the expected or capacity-weighted average is insufficient because a series path fails at its narrowest edge.

Substitution latency. Redundancy exists only when an alternative can enter before the service obligation expires. For each edge publish a substitution time $\ell_e(s)$ and the workload's

maximum tolerated interruption L_W. An alternative with $\ell_e(s) > L_W$ does not contribute to DVC in that scenario, even if it eventually restores capacity. Spare hardware in another jurisdiction is not current redundancy if export approval, installation, synchronization, or key rotation takes longer than the SLA permits.

Common-cause dependencies. Nominally separate facilities may share a transformer vendor, firmware signer, cloud identity layer, cable, gas market, cooling basin, stablecoin, legal safe harbor, or market maker [Doomberg 2026c]. Assign each edge one or more dependency groups g. A scenario can haircut or remove all edges in a group simultaneously. Geographic dispersion without dependency-group dispersion is not independent capacity.

Relationship to Sovereign Optionality. Sovereign Optionality is a capacity-weighted composite of substitutability attributes that *correlates with* independent pathway count but does not count pathways. DVC asks how much usable service reaches the sink after the scenario binds, and is the measure that actually enumerates flow. A network can have high optionality and low DVC because all alternatives are small or slow; it can have high benign DVC and low optionality because one large path supplies nearly everything. Report both, with the active binding constraints and dependency groups, rather than combining them into a reassuring scalar.

Issuance discipline. Work Credits, capacity vouchers, and any other forward service claim must be bounded by a conservative stress-adjusted quantity,

$$Q_W^{\text{issue}}(t) \;\le\; |t| \cdot \min_{s \in \mathcal{S}_{\text{constitutional}}} \text{DVC}(W, t, s) \cdot h_{W,t}, \qquad 0 < h_{W,t} < 1,$$

where $|t|$ is the length of the issuance interval, DVC is a rate, and $h_{W,t}$ is a published prudential haircut. The interval factor on the right is what makes a stock (claims issued) comparable to a rate (capacity delivered); it belongs in the bound, not in a gloss beneath it. $\mathcal{S}_{\text{constitutional}}$ is the pre-declared set of scenarios the constitution treats as binding for issuance. Scenarios carrying zero flow do not vanish from the minimum—a scenario in which nothing is delivered binds issuance at zero, which is the correct answer—and they are excluded only if the constitution pre-declares them as force-majeure. That exclusion list is itself a published governance parameter, subject to the same change-disclosure rules as any other; there is no default exclusion. FCR, FER, and PIDL artifacts supply evidence for edge capacities; they are not themselves capacity. Governance may not replace this bound with nameplate power, installed hardware, gross proof throughput, or a benign-state average.

A rate is not a joint completion. DVC is a steady-state flow of independently settled units. Work delegated across operators (Chapter 7) arrives as trees of units that must *all* complete

inside one refund window, and a tree that spans operators spans grids and jurisdictions, so the edge whose loss aborts the most leaves is a common-cause dependency in the sense above. A network can therefore deliver its DVC in units per interval and still fail most deep trees during a disruption. The joint-completion scenario of Section 22.9 measures that gap. It tightens no issuance bound in this version; it reports whether a bound stated in flow overstates what delegated demand can actually consume.

17.10 Resilience Versus Affordability: The Layer 0 Cost Tension

The previous section asks for resilient, dispersed, fuel-diverse, backup-equipped capacity. The thesis elsewhere insists that verification must stay cheap: "anyone can verify" is the hinge, and Section 30.6 treats its failure as fatal.

These two demands pull against each other, and the tension should be stated rather than glossed.

> Resilience is insurance. Insurance has a premium. A jurisdictionally dispersed, fuel-diverse, backup-equipped fleet costs more per verified unit than a single facility sitting on cheap interruptible power.

If hardening Layer 0 raises cost per verified unit, then hardening Layer 0 pushes the stack *toward* the very red line that hardening was supposed to protect. A thesis that demands both cheap verification and expensive redundancy without reconciling them has an open flank.

17.10.1 Two Different Exposures

The tension dissolves partially once two things are separated that "energy cost" conflates.

The verification side. Checking a proof happens on reference hardware, at the edge, in small amounts. Its physical exposure is not bulk industrial power. It is:

- **Reference hardware obtainability.** Export controls, sanctions, or platform lockdown can make the reference verifier unavailable, or available only in attested-and-permissioned form.
- **Edge power and connectivity cost** relative to local income, which is what makes verification affordable *in practice* rather than in principle.

The proving side. Generating proofs and running verified compute consumes bulk power. Its exposure is:

- **Energy price and firmness** at facility scale.
- **Competition for the same electrons** from industrial load and state-directed AI build-outs.

- **Curtailment, rationing, and interconnection denial**, which are policy instruments, not market outcomes.

Conflating these produces sloppy claims in both directions. Bulk electricity prices do not directly break "anyone can verify," because verification is not a bulk-power activity. What breaks the hinge on the verification side is loss of access to unprivileged reference hardware. What bulk energy conditions threaten is the *proving* economy: Work Credit issuance, geographic concentration, and enclosure.

Scope of "Exogenous"

Section 22.7 describes the Real-Resource VerifyPrice SLOs as "constitutional" and "exogenous to token price." That is correct and worth keeping: the SLOs must not be redefined because the market moved.

But exogenous to token price is not exogenous *simpliciter*. Real-resource VerifyPrice remains endogenous to the physical and jurisdictional conditions under which reference hardware and power are obtainable. Those conditions are partly set by states. The constitutional claim should be read with that scope, not as a claim that the metric floats free of politics.

17.10.2 Disruption-Adjusted VerifyPrice

For the proving side, comparing a resilient fleet against a brittle one at benign-state spot cost is the wrong comparison. It prices the premium and ignores what the premium buys.

What this subsection prices is therefore **fleet resilience of verification**: $\text{VerifyPrice}_s(W)$ below is the cost of verification for workload W in state s, where s runs over the five price states. Two same-named metrics must be kept apart. *Physical VerifyPrice* (Section 22.8) is the Layer 0 audit cost of checking FCR claims—a measurement of how expensive it is to verify the physical collateral story. The *VerifyPrice SLOs* of Layer 4 are the constitutional cost band on verifying a workload's receipts. They are different metrics that share a surname, and the decision rule at the end of this section reads on the latter.

The object under measurement is a *hardening fleet*: a collection of proving facilities deliberately built redundant, dispersed, and fuel-diverse so that verification survives the states in which it is most needed. Hardening is not decoration. A fleet that has paid for geographic and jurisdictional dispersion is one whose verification still runs when a single region is curtailed, a single transformer vendor is compromised, or a single jurisdiction is excluded—which is precisely when the cost of checking becomes the difference between a claim being auditable and a claim being whatever the surviving operator says it is.

Let s index physical disruption states—normal operation, regional curtailment, extended outage, hardware supply interruption, jurisdictional exclusion—with probabilities π_s, and let

$F_s(W)$ be the empirical distribution of *per-run verification cost* for workload W observed (or, for states not yet observed, modeled from receipts) in state s. The disruption-adjusted figure is computed from the **π-weighted pooled distribution**

$$F_{\text{adj}}(W) \;=\; \sum_s \pi_s \, F_s(W), \qquad \text{VerifyPrice}_{\text{adj}}(W) \;=\; \big(p_{50},\, p_{95}\big) \text{ of } F_{\text{adj}}(W),$$

that is, the mixture distribution in which a run is drawn from state s with probability π_s, and the percentiles are then read off the mixture. This matters because a π-weighted average of per-state p_{95} values is *not* a percentile of anything: percentiles do not average, and a fleet whose outage-state p_{95} is catastrophic would see that tail diluted linearly by a small π_s rather than showing up, as it should, in the pooled tail. Only the failure-rate field is a mean, and only it averages linearly: $\text{fail}_{\text{adj}} = \sum_s \pi_s \, \text{fail}_s$.

Two scope notes are recorded with the definition, and kept brief here. First: the states s are the **five price states** enumerated above, not DVC's capacity scenarios (Section 17.9)—two distinct enumerations serving distinct purposes, and they must not be merged. Second, on dimension: the pooling is done **separately for time and for cost**, yielding the five-field tuple $(p_{50,t},\, p_{95,t},\, p_{50,c},\, p_{95,c},\, \text{fail})$, the first four as pooled percentiles and the last as a linear average; when a scalar headline is quoted for $\text{VerifyPrice}_{\text{adj}}$ it is the $p_{95,c}$ field. (Physical-component tuples such as $(t_{50}, t_{95}, e, m, f)$ are views derived from receipts via a per-run cost model, and cost percentiles are percentiles of per-run cost.)

Disruption-Adjusted VerifyPrice

The percentiles of verification cost over the π-weighted pooled distribution of per-run costs across physical disruption states, rather than the observed cost under benign conditions. A convenience summary; no red line or acceptance rule reads on it.

A resilient fleet has higher $\text{VerifyPrice}_{\text{normal}}$ and lower $\text{VerifyPrice}_{\text{adj}}$ whenever the π-weighted saving in the disruption states outweighs the π-weighted premium in the benign state. Writing n for normal operation, let $\Delta_n \geq 0$ be the **premium**—the amount by which the resilient fleet's benign-state cost exceeds the brittle fleet's—and, for each $s \neq n$, let $\Delta_s \geq 0$ be the **saving**—the amount by which the resilient fleet's cost in state s falls below the brittle fleet's. Then the resilient fleet wins on the adjusted figure whenever $\sum_{s \neq n} \pi_s \Delta_s > \pi_n \Delta_n$. (For the pooled percentiles this holds as a first-order statement about location; the exact comparison is made on the pooled tuples themselves.) A brittle fleet optimised purely for spot cost reports an attractive headline number and carries the tail.

The π_s are estimates and should be published as such, with their basis, rather than presented as measurements. An operator who can move π_s by assertion can make any fleet look resilient; disruption-state probabilities are therefore governance parameters subject to the telemetry-capture constraints of Section 30.6. Because operators may set self-flattering π,

cross-fleet comparisons use a common reference vector π^{common} published by the observatory; an operator's self-published π may be shown alongside for context but is never used for cross-fleet ranking.

One further limit belongs beside the definition, because the mixture form invites a misreading it should prevent. Pooling fixes the arithmetic—the p_{95} of the mixture is a genuine percentile—but it does not fix the weighting: a state assigned $\pi_s < 0.05$ contributes fewer than five percent of the pooled runs, and its entire cost distribution can sit above the pooled p_{95} without moving it. Resilience is a property of the tail conditional on the disruption, not of the pooled tail: a fleet can hold a flattering $\text{VerifyPrice}_{\text{adj}}$ while its outage-state cost is catastrophic, if the assigned π_s for that state is small enough. The thesis's own treatment of capacity therefore declines expectation forms in the same situation — Delivered Verified Capacity is a scenario max-flow, with each scenario reported separately, rather than collapsed into a single weighted number (Section 17.9). $\text{VerifyPrice}_{\text{adj}}$ is accordingly a **convenience summary for comparison across fleets, not a quantity any red line or SLO reads on**. The load-bearing publication is the per-state vector $\{\text{VerifyPrice}_s(W)\}$ itself, published state by state alongside the π_s used, so that a skeptical reader can reweight with their own probabilities and see exactly where the summary flatters.

17.10.3 The Decision Rule

Priority ordering matters here, because "resilience" is exactly the kind of word that can be used to excuse blowing through a constitutional target.

> Accept a benign-state resilience premium only if **(i)** every pre-declared disruption state's VerifyPrice remains inside its per-state bound, **and (ii)** the pooled π_{common}-weighted $p_{95,c}$ falls, **and (iii)** benign-state Layer 4 VerifyPrice remains inside the constitutional SLO band.

The terms are defined so the rule can be applied rather than admired. The **per-state bound** in clause (i) is the constitutional Layer 4 VerifyPrice SLO band multiplied by a published **state multiplier** $m_s \geq 1$, with $m_n = 1$ for normal operation: verification in a curtailment or exclusion state is permitted to cost more than in the benign state, by a factor the constitution declares in advance and publishes beside the π_s, and no more. The multipliers are governance parameters under the same change-disclosure rules as the SLO band itself. Clause (iii) reads on the **Layer 4 VerifyPrice SLOs** only—the constitutional cost band on verifying workload receipts. It does not read on Physical VerifyPrice (Section 22.8), which is an audit-cost measurement with no SLO band; Physical VerifyPrice is published alongside as context and enters no clause of this rule. $\text{VerifyPrice}_{\text{adj}}$ appears only in clause (ii) and only as a **tie-breaker between fleets that have already passed (i) and (iii)**; a premium that improves the pooled figure

while pushing any disruption state past its bound is rejected by clause (i) regardless of what the pooled figure says.

The final clause is not negotiable. Resilience is subordinate to the hinge: a fleet so hardened that ordinary users can no longer afford to verify has defeated the purpose of hardening it. Redundancy that breaches the SLO band is not insurance; it is the failure mode wearing insurance as a costume.

This gives the thesis a definite position on "resilience over efficiency," worth stating plainly because the slogan travels widely and is usually asserted rather than priced:

- A **bounded, measurable** premium that buys a large reduction in tail risk is acceptable, and should be visible in FCR data rather than hidden in operator margin.
- An **unbounded** premium, or one that cannot be shown to reduce disruption-state cost, is not resilience. It is capital misallocation with a security narrative attached.

The FCR fields of Table 17.6, aggregated as the sovereign optionality index of Section 17.8, are precisely the instrument that distinguishes the two cases. Without them, "we are building resilience" is unfalsifiable. With them, it is a number an allocator can check.

Where this sits in the energy-security literature. The trade-off named here — redundancy and optionality priced against the cost of carrying idle capacity — is the central, long-standing question of energy systems reliability economics: the "resource adequacy" debate over how much excess capacity a grid should carry, and who pays for it. That literature (regulatory and academic, from NERC reliability standards [NERC annual] through the capacity-market literature) has spent decades converging on the same conclusion this section reaches by a shorter route: resilience that is not metered becomes rent. The thesis does not import the literature's models, and its FCR/DVC instruments are its own; but the position taken here is not contrarian, and the slogan it prices is the same slogan that literature prices. Analyst commentary used elsewhere in this Part (Doomberg, in Sources [Doomberg 2026a]) supplies the scenario framing; the regulatory economics supplies the skepticism.

17.11 Operational Patterns: Profiles, Upgrades, and Failure Modes

Layer 0 is not static; hardware evolves, breaks, and gets deprecated. It needs patterns for **living with that churn**.

17.11.1 Hardware Profiles and Workload Binding

Each canonical workload W in Layer 4 is associated with one or more **acceptable hardware profiles** $(H_1, H_2, \ldots)$. Profiles define:

- Minimum performance characteristics (to keep VerifyPrice in target bands).

- Acceptable side-channel leakage budgets.
- Attestation/sampling histories.
- Known caveats (e.g., "avoid profile H3 for workloads with secret inputs; leaks are too strong").

When PAL compiles a workload or the router assigns work, it can target specific profiles, diversify across profiles, or refuse high-sensitivity workloads to marginal profiles. This is how **Layer 0 informs the market** rather than hiding under it.

17.11.2 Upgrades and Deprecations

Hardware ages; bugs and backdoors are discovered; fabs change hands. Layer 0 needs clear **life-cycle rules**:

Onboarding: New profiles go through a probation period with extra sampling and conservative Work Credit weights.

Deprecation: When a profile is compromised or obsolete, higher layers stop accepting new Work Credits reserved on it, risk-flag existing credits, apply collateral haircuts, or impose prospective ineligibility, and publish an incident report at Layer 6.

Migration: Proof factories and corridor operators need technical paths to migrate workloads off deprecated profiles without massive downtime.

In monetary terms, this is the **hardware analog of a bond downgrade**: transparent, describable, and priced, rather than silently swept under the rug.

Non-Discretionary Downgrade Rules

To prevent "ex post discretionary default" critiques, deprecation and haircut decisions follow a **predictable severity framework**, not ad hoc governance.

Severity	Definition	Examples
S0 (Watch)	Potential issue; under investigation	Anomalous side-channel readings; unverified third-party report
S1 (Warning)	Confirmed issue with limited impact	RNG bias below 1%; isolated firmware bug
S2 (Critical)	Confirmed issue with systemic impact	Backdoor in >5% of sampled units; key extraction demonstrated
S3 (Emergency)	Active exploitation or catastrophic risk	Widespread key compromise; vendor collusion confirmed

Table 17.8: Severity levels for hardware profile issues.

Process guarantees:

1. **Evidence threshold:** S1 requires independent lab confirmation; S2 requires reproducible demonstration; S3 requires active exploitation observed.

2. **Multi-party decision:** Severity escalation requires sign-off from $\geq$3 of 5 designated hardware security reviewers.
3. **Appeal process:** Profile sponsors can challenge findings within 14 days.
4. **Sunset, not confiscation:** Even at S3, existing credits remain eligible at 75% of protocol-recognized collateral value, subject to the issuance bound of Section 17.9: if removing the profile's capacity from DVC leaves claims outstanding above the bound, the excess is impaired under the published retirement rules rather than carried. This is a statement about the protocol's own accounting, not a guarantee about market price: no protocol can promise what its asset trades for after a catastrophic hardware compromise, and this document makes no such promise.
5. **Transparency:** All severity determinations are published.

17.11.3 PQ and Cryptographic Agility

Some Layer-0 assumptions are about **cryptography**, not just silicon: signature schemes in ROMs, hash functions in hardware accelerators, and RNG primitives.

Layer 0 insists that hardware and firmware expose enough **configurability** to migrate to post-quantum or new primitives without throwing away entire fabs. Hardware profiles document cryptographic agility (e.g., "can switch hash from X to Y via firmware; signature scheme fixed").

17.12 Stress Tests for Layer 0

Every claim in Layer 0 should map to a testable stress scenario:

If a stress test fails, Layer 0 **degrades visibly**: profiles are flagged, Work Credits tied to them are risk-flagged and may face prospective ineligibility, and Layer 6 publishes an incident report.

17.13 What Layer 0 Exports to Higher Layers

Higher layers consume a **small set of artifacts and APIs**:

1. **Hardware profiles (HIDs).** Compact identifiers + dossiers describing profile properties, sampling history, and current status.
2. **Attestation receipts.** Machine-level statements binding device $\rightarrow$ profile, firmware $\rightarrow$ hash, measurement $\rightarrow$ nonce/time.
3. **Power and health telemetry.** Streams of power usage, failure rates, and uptime patterns.
4. **Incident and status flags.** Signals like "Profile H3 compromised; do not accept new work."

Claim	Stress Test
Profile H is honest within bounds X	Feed adversarial inputs; measure deviation from spec; decap sample devices
Sampling methodology detects tampering	Red-team: insert known-bad units into supply; measure detection rate
RNG entropy meets threshold	Run NIST SP 800-90B entropy-source estimators [Turan et al. 2018] on the raw noise source (min-entropy, not statistical randomness); run Dieharder/TestU01 on the conditioned output only; inject synthetic bias; verify detection
Power telemetry is honest	Cross-check metering with grid records; test resilience to meter spoofing
Migration from H1 to H2 works	Simulate H1 deprecation; measure latency and failure rate during migration
Capacity bounded by DVC can serve delegated work	Run the reference delegation tree of Section 22.9 across facilities under each disruption scenario; report completion probability inside the refund window and the dependency group whose loss aborts the most leaves

Table 17.9: Layer 0 stress tests.

Chapter 18

Layer 1: Reachability

All of this math still travels as packets.

Every proof, every private swap, every verified FLOP ultimately crosses a handful of cables and radios that a small number of operators can see and shape. A store of value that survives yield-curve control but dies when a few IXPs collude is not a store of value; it is an overlay.

Comms resistance is the stack's oxygen: the property that air still flows when someone leans on the hose.

Prior art this chapter stands on. The layer's components are deployments, not proposals, and the design borrows deliberately: BIP-324 encrypted P2P transport [Mehta et al. 2023] (shipped in Bitcoin Core 26.0, enabled by default in 27.0 [Bitcoin Core 2024]) is the wire-format template for encrypted peer transport; Tor [Dingledine et al. 2004], I2P, and the pluggable-transport lineage—obfs4, meek, Snowflake, and the refraction-networking line ending in Conjure [Frolov et al. 2019]—are the first-class transports and the camouflage fallbacks; BOLT12 Offers [NoBS Bitcoin 2024b] are the receiver-private rendezvous pattern; and the VerifyReach measurement program of Section 18.5 follows the OONI methodology [OONI 2012–] — volunteer vantage points, canary requests, ASN-stratified public aggregation — that has measured real-world censorship at scale since 2012. The thesis's contribution is not any of these mechanisms; it is binding them into a monetary stack whose falsifiability depends on them, and treating their telemetry as monetary instruments. Full entries in Sources.

18.1 Threat Model: What the Stack Must Survive

The comms threat model has four main faces:

Backbone controls. IP and prefix blocking, BGP blackholing, DNS and SNI filtering, DPI-based resets, QoS throttles on known P2P handshakes.

Exposure and linkage. Static addresses and reusable invoices let adversaries map who is getting paid; mempool and gossip surveillance reveal patterns.

Eclipse and routing capture. A node that only talks to a small set of peers in the same ASN can be effectively isolated.

Platform risk at the edge. App-store takedowns, CDN firewalls, corporate endpoint policies that classify P2P clients as malware.

18.2 Design Rules: Protocol Posture Under Pressure

1. **Transports must be encrypted by default, and must not mistake encryption for camouflage.** Clear-text, distinctive handshakes invite classification, and BIP-324-class v2 encrypted transport removes them. But a stream of uniform-random bytes is itself a fingerprint: Wu et al. (USENIX Security 2023) [Wu et al. 2023] documented that the Great Firewall has, since late 2021, blocked fully encrypted traffic heuristically—by byte-entropy and printable-ASCII tests on the first packet—without needing to identify the protocol at all. "Boring on the wire" therefore splits into two distinct properties, and the stack needs both: *encrypted* transports (BIP-324, obfs4) that defeat content inspection and protocol fingerprinting, and *look-like-something* transports that defeat entropy heuristics by resembling traffic the censor cannot afford to block. It is tempting to treat the first as sufficient. It is not.
2. **The stack must be transport-agile.** No single path (TCP/TLS, Tor, I2P, QUIC) should be a single point of failure, and the fallback ladder must include at least one look-like-something transport per Section 18.3.
3. **Receiver privacy must be the default.** Payment rendezvous should be "addressless"—BOLT12 Offers, path-blinding, and shielded pools.
4. **Settlement must be refund-safe under squeeze.** Atomic, adaptor-signature flows with clear timeout and refund semantics.
5. **Edge admission must remain open.** No "special" relays or whitelisted entry points.
6. **Comms health is part of the telemetry regime.** If reachability silently degrades, neutrality degrades with it.

18.3 Mechanisms: What Needs to Ship

Encrypted, multi-path P2P transports. Full nodes speak encrypted P2P by default (BIP-324-class, built on a Noise-style authenticated key exchange—Noise is a handshake framework and supplies the encryption, not the camouflage), with Tor and I2P as first-class citizens. Against censors that fingerprint protocols, obfs4 removes the recognizable handshake and is the encrypted fallback of first resort. Against censors that block on entropy, a second class of transport is needed, and it is first-class here rather than an afterthought: **TLS mimicry** via uTLS, so the client handshake is byte-identical to a mainstream browser's; **Snowflake**, which carries traffic inside WebRTC sessions indistinguishable from video calls, brokered through ephemeral volunteer proxies; **meek** and domain fronting, which ride inside TLS to a large CDN whose collateral cost of blocking is high; and **refraction networking** (Conjure) [Frolov

et al. 2019], in which a cooperating transit ISP redirects flows addressed to unused IP space in its range, so the censor sees a connection to a decoy and blocking it means blocking the ISP. Each of these has been defeated somewhere; none has been defeated everywhere at once, and the reachability program of Section 18.5 exists to say which one is working where. Relays in high-interdiction ASNs advertise the transports they accept, and the client's fallback ladder is ordered by measured success in its region rather than by preference.

Transaction-origin privacy. The threat model of Section 18.1 names mempool and gossip surveillance, and encryption on the wire does nothing about it: a well-connected observer who receives a transaction first from node A learns that A originated or was near the origin, encrypted link or not. Transaction relay therefore uses **Dandelion++-style** two-phase propagation [Fanti et al. 2018]—an anonymity phase in which the transaction is forwarded along a random privacy graph with no fan-out, followed by a diffusion phase from a node that is not the origin—so that the first-spy estimator's advantage is bounded rather than total. Layer 1 owns the relay graph; whether the transaction's contents are also private is a Layer 5 question.

Receiver-private routing for payments. BOLT12 Offers [NoBS Bitcoin 2024b]: receiver-private, reusable invoices. Two distinct mechanisms are in play and must not be conflated. **Onion routing** hides the route from the intermediaries: each hop learns only its predecessor and successor. **Blinded paths** (BOLT12) hide the receiver from the sender: the receiver constructs the last hops of the route and encrypts them so the sender can pay without learning the receiver's node identity. The first protects both parties from the network; the second protects the receiver from the payer. Federated ecash [Fedimint Project 2024–25] and shielded pools provide **endpoint firebreaks**—points at which the payment graph is cut and cannot be followed through—with one residual to state plainly: a federated ecash mint is custodial. The federation holds the funds, and the firebreak is bought by trusting a threshold of its members not to steal or freeze them. It is a privacy tool with a custody cost, not a non-custodial one, and users choosing it are choosing that trade.

Settlement survivability. Adaptor-signature atomic swaps as the default for cross-asset payouts. "Abort & refund" as a first-class, well-signposted action.

Topology and anti-eclipse hardening. Peer-set diversity and rotation; connections spread across geographies and ASNs; gossip protocols favor multiplicity of paths.

18.4 How Layer 1 Anchors the Triad

Communications resilience is not an orthogonal concern; it is the medium in which the triad either lives or suffocates.

For Privacy, encrypted transports and receiver-private endpoints are what keep the network itself from nullifying ledger-level secrecy. Shielded pools and privacy coins are only as private as their membrane to the outside world. If every shielded payout can be traced to a stable IP graph and a static address, "lawful privacy" collapses into a thin veneer over a rich flow-of-funds analysis. Layer 1's job is to ensure that settlement paths are as hard to pin down as the flows they carry: private not only in state, but in motion.

For Proofs, communications resilience keeps verification a public act rather than a priestly privilege. A world where only a small set of well-positioned nodes can fetch and check proofs is a world where VerifyPrice has quietly become an internal KPI rather than a public commodity. If any honest machine with modest connectivity can still reach a verifier over at least one path, proof markets and receipt ledgers remain subject to universal scrutiny. When reachability fragments, the "public" in public verification becomes aspirational.

For Compute, Layer 1 ensures that useful-work mining and verified inference do not devolve into "whoever still has a clear line to the router wins." Proof-of-Useful-Work schemes depend on open admission: a wide, geographically and topologically diverse set of provers and miners competing to satisfy claims. If adversaries can choke ingress to a few ASNs or clouds, PoUW degenerates into a club good. Comms health metrics and open-admission design at the edge keep compute supply neutral and keep the token's claim ("backed by globally demanded work") from turning into "backed by whichever datacenter the regulator likes."

All of this math still travels as packets. Layer 0 keeps the machines honest and powered; Layer 1 keeps them in conversation when it is no longer convenient for them to be. The triad's monetary ambitions depend on both. Without verifiable machines, nobody knows what happened. Without resilient communications, nobody knows it in time, or at all.

Residual risks: Encrypted transport and obfuscated routing are ingredients, not complete answers to censorship. Residual threats include traffic analysis (pattern and timing correlation), active probing (protocol fingerprinting despite camouflage), endpoint seizure (physical confiscation of devices), long-term sustained shutdowns (where no transport can help), and legal coercion at the social layer. Layer 1 reduces these surfaces; it does not eliminate them. The honest claim is degradation resistance, not censorship immunity.

18.5 VerifyReach: Communications Telemetry

Layer 1 introduces **VerifyReach** as the communications analogue of VerifyPrice:

- **Reachability metrics**: fraction of vantage points from which key services are reachable.
- **Degradation patterns**: which networks experience blocking or throttling.
- **Transport diversity**: percentage of traffic over each transport class.

Target SLOs:

- Core infrastructure reachable from $\geq 95\%$ of sampled ASNs in **uncensored regions**, where "uncensored" means outside the pre-declared list of known-censored regions maintained by the observatory (the "Censorship regimes" row of the sampling frame below). That list is a published governance parameter; a region moves onto it by measured, sustained filtering, not by assertion.
- $\geq 70\%$ reachability for **known-censored regions**, measured over the ASNs sampled inside those regions.
- p95 time-to-first-connection < 10s uncensored; < 30s censored.
- In **uncensored regions**, no country-level view sees more than 10% persistent reachability degradation without triggering incident handling. In known-censored regions, where 30% degradation is already within the SLO, this trigger would fire permanently and mean nothing; there, incident handling is triggered instead by a *fall* of ≥ 10 percentage points from that country's trailing 90-day reachability, so the rule reads on change rather than on a level the region never met.

The denominators are stated because the falsification test of Section 20.12 reads on them: that test counts ASNs in uncensored regions, not globally, so that the 70% censored-region SLO and a $> 20\%$-unreachable failure condition are not in contradiction. A region can be at 65% reachability and be an incident; it cannot, by itself, falsify the thesis.

VerifyReach feeds into Layer 6 governance: if reachability collapses in a region, incident response kicks in (alternative transports are promoted, routing is adjusted, and the degradation is visible on public dashboards).

18.5.1 VerifyReach Measurement Specification

Like VerifyPrice, VerifyReach requires a rigorous measurement methodology to prevent gaming and ensure credibility.

Sampling Frame:

Vantage Points:

Dimension	Minimum Coverage	Rationale
ASNs	≥ 500 distinct ASNs globally	ISP-level blocking
Regions	$\geq$ 50 countries; $\geq$ 10 per inhabited continent (five-continent convention)	Regional censorship
Censorship regimes	Coverage of known filtering states (CN, IR, RU, etc.)	Worst-case reachability
Network types	Residential, mobile, enterprise, datacenter	Filtering varies by type

Table 18.1: VerifyReach sampling frame requirements.

- **Community nodes:** Volunteer-run measurement agents (similar to OONI probes [OONI 2012–]), paid small amounts from the assurance budget (Section 25.6) — not in Work Credits, which mint only against registry workloads.
- **Independent labs:** At least 3 organizations (academic, NGO, commercial) run measurement infrastructure. No single operator controls $> 30\%$ of vantage points.
- **Diversity requirement:** Vantage points must span $\geq$ 20 countries and $\geq$ 100 ASNs for measurements to be considered valid.

The two threshold levels do different jobs and are not redundant. The sampling *frame* targets ≥500 ASNs and ≥50 countries; the *validity floor* is the minimum below which measurements are not published at all (≥100 ASNs, ≥20 countries). Continent arithmetic follows the five-continent convention used throughout this document (≥10 per inhabited continent).

Metric	Definition	Target SLO
succ_1(N,R)	Fraction reaching service N via primary transport within 30s	$\geq 95\%$ / $\geq 70\%$
succ_2(N,R)	Fraction reaching via any transport (incl. fallbacks) within 60s	$\geq 99\%$ / $\geq 85\%$
ttfc(N,R)	Time-to-first-connection (p50, p95), *conditional on successful connection*	p95 $<$ 10s / $<$ 30s
failure_class	Taxonomy: DNS, TCP RST, TLS, timeout, active probe	Published per region

Table 18.2: VerifyReach metrics and targets. Targets shown as uncensored / known-censored.

Metrics: The ttfc and succ targets bind **both**, because they measure different things: ttfc percentiles are conditional on successful connection and say nothing about how often connection fails, while succ_1/succ_2 measure the unconditional success rate and say nothing about how slow the successes were. Neither subsumes the other.

Adversarial Robustness:

Publication:

- Raw measurement data (anonymized to protect vantage point operators) published daily.
- Aggregated dashboards updated hourly.
- Quarterly reports summarizing regional trends, incidents, and transport effectiveness.

Threat	Mitigation
Spoofed vantage points	Periodic challenges (fetch and sign specific data); anomalous behavior triggers exclusion
Selective treatment	"Canary" requests that should succeed; if canaries fail but targets succeed, point is flagged
Measurement capture	Cross-check results from different operators; divergence > 10 percentage points (measured in absolute percentage points between independently computed aggregates) triggers investigation
Temporal gaming	Continuous measurements (not snapshots); 24-hour rolling averages published

Table 18.3: VerifyReach adversarial robustness measures.

- All measurement code is open-source and reproducible.

Monetary Consequences: If VerifyReach for a region falls below thresholds:

- WC sold by operators in that region may face **regional risk premiums** (higher collateral requirements).
- Corridors primarily serving that region are flagged; users see warnings before transacting.
- Incident response is triggered: alternative transports promoted, routing adjusted.

This makes VerifyReach not just a dashboard metric, but an **input to economic risk pricing**.

Chapter 19

Layer 2: Distribution & Execution

Software distribution is where lofty protocol guarantees meet the boring reality of phones, laptops, and routers that have to run code. In a repression cycle, app stores are throttled, DNS is poisoned, and "safety" policies become chokepoints.

It is far easier to lean on an app store, certificate authority, or CDN than to break a new primitive. If privacy, proofs, and compute are to support an associated base-asset monetary candidate, then the binaries that implement those services must keep moving even when networks are hostile.

Prior art this chapter stands on. The supply-chain invariants below (I1–I5) are not novel claims; they are the TUF/Uptane guarantee set [TUF 2011–] [Kuppusamy et al. 2016] — threshold signing roles, freshness and anti-rollback, reproducibility, compromise response, multi-path delivery — restated for this stack, and the transparency-log mechanism is Certificate Transparency [Laurie et al. 2013] applied to release artifacts, with Reproducible Builds [Reproducible Builds 2014–] as the verification floor. Layer 2's contribution is narrower: adapting these to a monetary client whose update path must survive exactly the states in which the asset's holder-side service is supposed to deliver, and wiring the update channel's health into VerifyReach telemetry. Entries in Sources under "Software Update Security & Transparency."

19.1 Threat Model & Objectives

An adversary's repertoire at this layer:

- App-store removals and policy bans.
- DNS poisoning, SNI and IP blocking, TLS interception.
- BGP hijacks or route blackholing for update servers.
- Certificate revocations; captive portals that substitute their own "secure updates."
- Targeted developer coercion and "emergency directives."

Objectives:

- **Reachability**: users can discover and fetch releases despite interference.
- **Authenticity**: users can verify what they fetched without trusting a platform.

- **Safety**: upgrades are atomic, reversible, and survive power loss.
- **Neutrality**: no single jurisdiction or vendor can gate updates.
- **Telemetry without doxxing**: reliability is measurable without surveillance.

19.2 Design Rules

Authenticity before reachability. Releases ship with threshold signatures (3-of-5) from independent maintainers, plus inclusion in an append-only transparency log.
Reproducible builds as default. At least two independent builders reproduce binaries from source; the build process emits deterministic artifact hashes and proofs.
Threshold keys and real revocation. Root keys live offline and are sharded; release keys rotate regularly; revocation is a signed artifact in the transparency log.
Binary transparency anchored to a proof ledger. The update log behaves like certificate transparency [Laurie et al. 2013] for releases, periodically anchored to a neutral proof ledger.
Atomic updates with safe rollback. Clients update into an inactive slot, flip a pointer only after verification, and revert automatically on failure.
Graceful degradation. Offline packages (USB, SD, QR) use the same receipts and keys.

Supply-Chain Invariants (Must Hold)

I1. Authenticity: Every release requires threshold signatures ($\geq$3-of-5) from independent maintainers in $\geq$2 jurisdictions, plus transparency log inclusion.
I2. Freshness: Clients enforce monotonic version counters *and* expiry-bounded timestamp and snapshot metadata, per TUF [TUF 2011–]. Version counters stop rollback (serving an older, signed release); they do nothing against a *freeze* attack, in which a mirror or a captured network path serves the current release indefinitely while a patched one exists. Timestamp metadata is therefore re-signed at short interval and carries an expiry; clients reject any metadata older than T hours (reference: $T = 24$, aligned with I4 and I5) and treat an expired timestamp as a failed update, not a successful no-op.
I3. Reproducibility: At least 2 independent builders reproduce each release. Divergence is release-blocking.
I4. Compromise response: If a signing key is compromised: revocation within 24 hours, affected releases quarantined, post-mortem within 14 days.
I5. Multi-path availability: Releases available via $\geq$3 independent channels. No single channel's failure blocks updates for $>$24 hours.

19.3 Reference Distribution Architecture

The architecture is deliberately multi-homed. The design assumes any single channel can become a chokepoint.

Control plane (discovery and metadata):

- Standard HTTPS endpoints behind diverse DNS providers.
- Tor onion services exposing the same manifests.
- Signed release announcements gossiped over generic protocols (Matrix, Nostr).
- Append-only transparency log with roots checkable by light clients, and **witnessed checkpoints**: checking an inclusion proof against a root proves only that the release is in *a* log the client was shown, and a log operator who serves different roots to different clients (a split view) defeats it without forging anything. Log checkpoints are therefore co-signed by an independent witness set—the pattern of the Sigstore/Rekor witness network and Go's checksum database—and a client accepts a root only when it carries a threshold of witness signatures, each attesting that the root is consistent with every earlier root that witness has seen.

Data plane (artifact delivery):

- Traditional CDNs and anycast mirrors in multiple jurisdictions.
- P2P swarms (IPFS-like or BitTorrent-class) using content addressing.
- Community micro-mirrors with signed manifests.
- Offline channels: USB/SD bundles, QR-encoded update chunks, radio relays.

All of these are interchangeable from the client's perspective. Whoever delivers the bytes fastest wins, but no one is trusted beyond the hash: the client verifies the final digest and only then installs.

19.4 Key Management & Release Process

Key management is where many otherwise sophisticated systems quietly re-centralize. A single HSM in a single jurisdiction becomes a soft power lever: compromise the humans or the box, and you own the distribution pipeline.

A capture-resistant release process needs:

- **Public release ceremony.** Changes land on protected branches; CI produces deterministic binaries; independent builders confirm bit-for-bit equality. A quorum of maintainers signs the targets manifest.
- **Emergency freeze and unfreeze.** Any maintainer can propose a freeze, but activation requires a threshold of signatures. Unfreeze requires an explicit, signed artifact explaining the

incident.

- **Compromise playbook.** Detect, revoke, rotate, and re-sign last-known-good releases. All steps produce receipts that anyone can audit.
- **Jurisdictional diversity.** Key shards and signers are distributed across legal zones, making it difficult for one government to unilaterally subvert the process.

19.5 Update Economics & Neutrality

Distribution has an economics layer just as proofs and compute do. Instead of pretending bandwidth and storage are free, the stack lets mirrors be market actors: they meter what they serve, emit receipts, and get paid for honest work.

- **Receipts for bytes**: mirrors export metered receipts signed by their nodes; a broker aggregates and pays out via privacy rails with SLA escrow.
- **Open admission**: any party can become a mirror by publishing bandwidth and uptime commitments and staking a small bond.
- **Anti-capture telemetry**: publish house-share, top-N mirror share, geographic/ASN diversity. If a CDN dominates beyond thresholds, route weight decays automatically.

19.6 Fallback Playbooks

Having explicit, rehearsed playbooks keeps censorship events from becoming existential crises:
Soft network blocks (DNS/SNI/IP): Bias clients toward alternative control-plane channels (DoH/DoT for resolution, Encrypted Client Hello so the SNI itself is not readable on the wire, onions), and toward content-addressed P2P for artifacts.
Hard regional blocks: Rely on offline kits (USB/SD bundles, QR sets) seeded through civic institutions and NGOs.
App-store bans: Publish signed sideload bundles with clear, localized installation instructions.
Key compromise: Hit the freeze button; revoke; rotate; and re-sign last-known-good releases, with all steps recorded in the transparency log.

19.7 Execution Integrity

The chapter's title promises execution as well as distribution, and delivering correct bytes is not the same as executing them correctly. Three requirements close the gap, each reusing machinery already specified rather than adding a trust root.
Reproducibility verified at install. I3 has two independent builders reproduce each release

before it ships. The client repeats the check on its own machine: at install it recomputes the artifact digest and matches it against both the threshold-signed targets manifest and the builder digests recorded in the transparency log, refusing to activate a binary on which those disagree. A release that reproduced at the ceremony but not on the client's path is a supply-chain incident.

Least privilege and sandboxing. The reference verifier runs with no network access it does not need, no filesystem access beyond its receipt store and keys, and no ambient authority over wallet or host, inside the platform sandbox (seccomp/Landlock, App Sandbox, or a WASM runtime). A compromised dependency may lie about one receipt; it must not be able to exfiltrate keys.

Runtime attestation through the same log. Where the host offers measured boot or a TEE, the running verifier's measurement is compared against the release digest in the transparency log; a mismatch is surfaced to the user and to telemetry as a signature-failure event. The log that proved the bytes were the released bytes proves the process is running them. Its limit is that of Section 17.5.3: the measurement is only as trustworthy as the attesting root of trust, and hosts without one degrade to the install-time check, reported as the client's execution-integrity grade.

None of this is novel; all of it is routinely omitted. The title is earned only if the client that verifies proofs is itself verified to be the client that was released.

19.8 Lawful Privacy in Distribution

The same architectural moves that make distribution censorship-resistant also make it more legible to good-faith auditors. Transparent logs, reproducible builds, and signed receipts give regulators something objective to look at:

- Which releases were shipped, and when.
- Which vulnerabilities were patched, and how quickly.
- Whether there were jurisdiction-specific forks or backdoors.

Crucially, auditors can examine this evidence without sitting in the middle of every update flow. The system does not ask to be exempt from scrutiny; it insists that scrutiny be applied at the level of receipts and logs rather than taps and implants.

19.9 Update Telemetry

To make distribution part of the triad's credibility, networks should publish continuously:

- **Update health**: p50/p95 metadata and artifact fetch times; per-path success rates; rollback and signature-failure rates.
- **Transparency evidence**: current log roots; inclusion proofs for each release.

- **Decentralization metrics**: top-N mirror share; geographic and ASN diversity.
- **Key material and incidents**: root fingerprints; release-key roster; revocation lists; incident reports.

Chapter 20

Layer 3: Identity & Claims

Identity in a repression-prone, AI-saturated internet cannot be "a file with your name on it." It must be capabilities you can prove on demand (age, jurisdiction, uniqueness, solvency, model ownership) without handing over the dossier.

Reputation should be the trail of receipts from prior correct behavior, not a database of personally identifying facts.

20.1 Why Identity Must Decouple from Doxxing

The web's trust default has flipped: seeing is no longer believing, and platforms strip provenance labels inconsistently [Harwell & Patton 2025]. Accounts and government IDs are brittle anchors; they leak power to gatekeepers.

What scales instead is **mechanism over memory**: claims backed by cryptographic evidence any honest party can check cheaply and without permission. This is the same shift the thesis makes for money—verification, not authority, is the arbiter.

20.2 Design Goals

Identity and reputation in this stack satisfy six goals:

1. **Privacy by default, proof by construction.** No standing, globally queryable identity graphs. All predicates are proven as needed, most often with zero-knowledge proofs.
2. **Selective disclosure and unlinkability.** Verifiers learn only the outcome of a predicate, not your name or the rest of your credential bundle.
3. **Cheap, public verification.** Any honest laptop can check evidence quickly, measured the same way proofs are measured with VerifyPrice.
4. **Open admission for issuers and verifiers.** Anyone can issue attestations bound to their own reputation and slashing.
5. **Hardware honesty for machine identity.** Machine claims are anchored in verifiable machines with open designs and sampled supply chains.
6. **Lawful privacy.** Regulators get provable, epoch-scoped visibility via receipts and per-

epoch disclosure keys (Section 23.7), not through custodial chokepoints.

20.2.1 Issuer Pluralism Without Issuer Anarchy

"Open admission for issuers" (goal 4) sounds like chaos: anyone can issue AML credentials, so criminals just issue their own. The resolution is that **issuer sets are policy objects**, not global permissions.

Policy Context	Example Predicate	Issuer Set Definition
AML-screened	"User is not on OFAC/EU sanctions lists"	Threshold of $\geq$2 from {Exchange A, Exchange B, Compliance Provider C, Regulator D}
Age-verified	"User is 18+"	Any licensed identity provider in user's jurisdiction, or threshold of $\geq$2 from global providers
Accredited investor	"User meets SEC accredited investor definition"	Licensed broker-dealer or accredited investor verification platform
Device integrity	"Device is L0-B or higher"	Attestation from the device's own TEE *plus* one of: an independent root-of-trust component on the same device (open-design RoT, Section 17.4), or a membership credential issued to the device's hardware profile by the sampling program (Section 17.5.4). A single device does not carry two vendors' TEEs, so "two TEE vendors" is not a requirement one device can meet

Table 20.1: Issuer sets by policy context.

Issuer sets by policy context:

How issuers enter/exit sets:

1. **Stake:** Issuers post the base asset as collateral under a defined rule: a fixed *per-context floor* set by the policy context's governance, plus a *per-credential increment* for every credential active (unexpired and unrevoked) in that context, so that the required stake is $S_{\text{floor}} + \delta \cdot N_{\text{active}}$ and is recomputed at each epoch from the issuer's own published credential count. "Proportional to the value of credentials issued" is not a rule, because that value is not computable; a count of active credentials is. This collateral role is one value-accrual mechanism, not proof of monetary premium. If credentials are later found fraudulent, stake is slashed.
2. **Audit receipts:** Issuers must publish periodic audit receipts proving their verification processes meet policy requirements. Audits are conducted by independent third parties.
3. **Slashing:** If an issuer's credentials are shown to be systematically false (e.g., issuing "not on sanctions list" to sanctioned entities), their stake is slashed and they are removed from

relevant issuer sets.

4. **Jurisdictional diversity:** For high-stakes predicates (AML, accredited investor), issuer sets require representation from $\geq$2 jurisdictions to prevent single-government capture.

How verifiers select policy bundles:

- Verifiers declare which policy contexts they require (e.g., "AML_screened AND Age_verified").
- The protocol resolves these to issuer sets.
- Presentations must include credentials from issuers in the relevant sets.
- Verifiers can add custom issuer requirements (e.g., "must include credential from Regulator D for EU users").

Why this is not centralization:

- No single issuer is required for any predicate.
- New issuers can join sets by meeting stake and audit requirements.
- Issuer sets are transparent and auditable.
- Users can choose which issuers to use within the set.

This model preserves open admission while ensuring that predicates in regulated contexts are backed by accountable issuers.

20.3 The Identity Kernel

Layer 3 treats identity as a claim machine rather than a name registry.
Identifiers are cheap, ephemeral, and scoped—pairwise keys or DIDs, not global user IDs.
Credentials are issued by many parties (exchanges, employers, DAOs, regulators, devices). Each is a capability statement, not a dossier.
Presentations are one-off proofs of predicates over credentials: membership, attributes in a range, conjunctions of statements. Typically zero-knowledge.
Receipts are the portable artifacts. Every accepted presentation produces a PIDL receipt binding the claim, proof hash, SLA tier, and timestamps.

This kernel compiles directly into the developer surfaces: PAL SDK to express "prove this predicate," PIDL receipts to carry results, and privacy corridors to pay without doxxing.

20.4 Human Identity: Personhood and Compliance Without Dossiers

What to prove, privately:

- **Uniqueness / Sybil resistance.** Rate-limited nullifiers, stake-and-slash, or attested-liveness ceremonies.

- **Eligibility.** Age, residency, licensing, or "member-of a KYC-screened set," proved with set-membership or range proofs.
- **Standing & solvency.** Proofs of reserves/income ranges, revealed under viewing keys during credit underwriting.

How it composes with the stack: A wallet invokes PAL to compile predicates to proofs, pays over privacy rails, and gets a PIDL receipt the counterparty can verify on-chain or in a browser in < 5s p95.

20.4.1 Sybil Resistance Menu

Sybil resistance (preventing one entity from creating unlimited fake identities) is a core challenge. Different mechanisms offer different tradeoffs; the stack supports a menu of approaches rather than mandating one.

Tradeoff guidance:

- **For maximum privacy:** Rate-limited nullifiers + economic PoW. No linkage, but limited action rate.
- **For regulated contexts:** Stake-based + issuer-set credentials. Some linkage to stake address, but accountability.
- **For machine identity:** Device attestation bounded by L0 grade. Accept fingerprinting risk for hardware honesty.
- **For proof of personhood:** Social ceremonies with ZK proofs. Higher coordination cost, but strong Sybil resistance.

Composability: Mechanisms can be combined. A high-stakes action might require:

1. Rate-limited nullifier (prevents spam),
2. Stake (ensures accountability),
3. Issuer credential (meets regulatory requirement).

Each layer adds cost and reduces privacy; the protocol allows verifiers to specify requirements and users to choose compliant paths.

20.5 Machine Identity: Capability, Not Brand

Machines (provers, miners, sensors, model endpoints) need identity, but it should look nothing like a cloud vendor account. Their "identity" is a statement of capability on a particular hardware profile coupled with receipts for past behavior.

When a device participates, it attests to what it is and what it did. Attestations link to open hardware profiles. Those attestations are then wrapped in succinct proofs so any chain

Mechanism	How It Works	VerifyPrice	Unlinkability	Best For
Rate-limited nullifiers	Each identity can perform N actions per epoch; nullifiers prevent reuse without linking	Low	High	Spam prevention, voting, airdrops
Stake-based	Identity actions require WC collateral; bad behavior $\rightarrow$ slashing	Low	Medium	High-stakes actions, market making
Social/liveness ceremonies	Periodic video calls, in-person events, or social graph vouching	Medium	Low-Medium	Proof of personhood, high-value credentials
Device-based attestation	Hardware attestation limits identities per device	Low	Low	Machine identity; bounded human use
Economic proof-of-work	Solving computational puzzles to create identity	Low (prover cost is high; the check is one hash)	High	Bot prevention, rate limiting

can verify them cheaply—and cheaply is all the wrapping buys. The proof establishes that a signature under the vendor's attestation key exists; the vendor key remains the root of trust for what the attestation says (Section 17.5.3). The trust a verifier places in "a device of profile P" rests on the sampling record and openness grade of P, not on the proof around its quote.

This is **reputation for robots**: SLA-backed capacity that accrues a track record via receipts, not by vendor logo.

For a software agent, capability is half of identity; the other half is a *principal binding*: a predicate showing which root can revoke or cap the agent, without naming whoever holds it (Section 7.6). Binding an agent to a firm or an operator needs nothing beyond the credentials above. Binding it to a natural person additionally needs the personhood predicate of Section 20.4. Where that predicate cannot be had without a real-name anchor, the "human with standing" class of Section 26.8.9 is reported as firm or unknown, and the Human-Revocable Share is not computable for those holdings.

20.6 Reputation: Priced Behavior, Not Personal Data

- **Positive reputation** accumulates as fulfilled obligations with proofs: deliveries within SLA, honest inferences, timely settlements, clean audits.
- **Negative reputation** accrues as rejects and slashes: failed audits, invalid proofs, settlement timeouts.
- **Scope and decay** are explicit: reputation is contextual and decays unless renewed.

Because receipts are portable and verifiable, reputation becomes a market primitive. A prover can advertise: "over the last 30 days I delivered PROOF_2^{20} within p95 $<$ 1s at cost $<$ \$0.001 with 0.04% failure"—and anyone can check.

20.7 Patterns That Keep Identity Private and Usable

zK-KYC / zK-AML. Users prove membership in a regulated allow-list without revealing which issuer or which entry.

Rate-limited actions. Per-relationship nullifiers cap spam without linking handles.

Credit without doxxing. A lender demands a proof set plus a bond; the loan contract auto-slashes on missed proofs. No one learns the borrower's legal name unless escrow triggers.

Creator & data provenance. Cameras on open hardware sign origin. The signed manifest itself does not survive distribution: C2PA manifests [C2PA spec] travel as file metadata, and most platforms strip metadata on upload—the Washington Post's 2025 test [Harwell & Patton 2025] found seven of eight platforms did. What survives is a **soft binding**: a perceptual hash of the content, registered together with the origin proof in the receipt ledger at capture time, so that a stripped copy can be matched back to its receipt by content rather than by

metadata. The residual is the analog hole and its digital equivalents—re-encoding, cropping, screen capture—which degrade or defeat perceptual matching; the registry lowers the cost of proving provenance for content that has been stripped, and does nothing for content that has been re-photographed.

Work credentials for models. Model stewards prove they own specific weights and that inferences came from those weights under defined parameters. The scope of "prove" is bounded by the state of the art: zero-knowledge proofs of inference are practical for small models and for individual layers or sampled operations of larger ones, and **frontier-scale ZK inference is a research problem** [Peng et al. 2025], not a mechanism this stack can specify. For large models the pattern is therefore **attest-and-sample**: the steward commits to weights, the serving hardware attests the model hash it loaded (subject to the attestation limits of Section 17.5.3), and a sampled fraction of inferences is re-executed or algebraically checked by an independent party, with the sampling rate and coverage published as the credential's assurance level.

20.8 Failure Modes & Guardrails

Any identity layer that matters will eventually be attacked. What distinguishes a durable architecture is whether failures are observable, bounded, and repairable without appealing to a central authority.

Centralized attesters. If one vendor or agency controls machine attestation or anon-cred issuance, identity becomes permissioned by decree. **Countermeasure:** Pluralism—open hardware profiles and lot sampling, which are the two things that actually reduce dependence on the attester, with ZK-wrapped quotes as the cheap transport for their results rather than a substitute for them (Section 17.5.3); multiple issuers whose credentials are interchangeable or combined via threshold logic.

Proof cost creep. If p95 verify time drifts up, only gateways can afford to check receipts. **Countermeasure:** Treat identity workloads as first-class citizens of VerifyPrice; treat regressions as Sev-1 incidents.

Bridge and settlement doxxing. If corridors leak metadata, identity collapses into routing tables. **Countermeasure:** Cryptographic atomicity (adaptor signatures), refund safety, and public settlement telemetry.

Reputation blacklists. Global "bad lists" are honey traps for political pressure. **Countermeasure:** Scoped, revocable policies tied to evidence; actors retain the ability to present counter-receipts.

20.9 How Identity Fits the Modular Stack

Identity and reputation are just another set of workloads in the same loop:

1. **Create / Compute.** A human or machine emits a claim: "I am unique and 18+," "this device captured this video," "this miner computed MATMUL_4096 under profile P."
2. **Prove.** The PAL SDK compiles the predicate into a circuit or useful-work audit, routes it to a prover market, and returns a PIDL receipt. "Circuit" is literal for identity predicates and small models; for large-model inference the compiled object is an attest-and-sample audit (Section 20.7), because a full ZK circuit for a frontier model is not something PAL can emit today.
3. **Settle.** The privacy rails kit executes refund-safe payouts across BTC↔XMR corridors (BTC↔shielded-ZEC when admissible) or shielded pools.
4. **Verify.** Anyone runs `verify(receipt)` locally, on-chain, or inside another proof.
5. **Telemetry.** VerifyPrice, anonymity-set size, swap success rates, and hardware-profile distributions are public.

The stack stops being "crypto plus KYC" and becomes one fabric where work, truth, and permission all pass through the same narrow gate: proofs anyone can check.

20.10 Identity/Reputation SLOs

For operators and allocators, identity becomes legible through metrics:

- **Predicate VerifyPrice.** The canonical five-field tuple $(p_{50,t}, p_{95,t}, p_{50,c}, p_{95,c}, \text{fail})$ per standard identity predicate, i.e. median and tail verify time and cost plus failure rate. Target: p95 under a few seconds, sub-cent cost.
- **Unlinkability.** Rate at which presentations can be linked by passive adversaries.
- **Settlement privacy.** Corridor success rates and anonymity-set sizes for identity-linked flows.
- **Hardware honesty.** Share of machine-signed receipts tied to open profiles.
- **Reputation health.** Fraction of decisions based on receipts vs. account flags.

Identity $\neq$ name; it is a set of provable capabilities. **Reputation $\neq$ biography**; it is a ledger of receipts. Bind both to privacy by default, publish VerifyPrice so anyone can check them cheaply, and anchor machine attestations in verifiable hardware.

20.11 Implementation Sketch

All of the above shows up in three simple patterns that developers wire into their apps:

Human $\rightarrow$ compliance without names. A wallet encodes a claim like: prove(age $\geq$ 18) $\wedge$ prove(residency $\in$ EU) $\wedge$ prove(member_of AML_screened_set). It sends this to PAL, which compiles the predicates, sends them to a prover market, and returns a PIDL receipt. The merchant verifies in $<$5s p95 without ever seeing the passport.

Machine → capability without brand. A sensor on hardware profile `open_camera_v1` captures video. The device emits an attested capture, PAL wraps it in a succinct proof, producing a provenance proof and receipt. The camera's "identity" is simply "a device of profile P that has produced N correct receipts."

Reputation → receipts, not profiles. A prover's handle accumulates receipts. When it advertises itself, it publishes: "Last 30 days: delivered PROOF_2^{20} under SLA Gold with p95 verify-time 0.8s, 99.9% success, zero slashes." Anyone can pull the underlying receipts and run the reference verifier.

20.12 Part III Summary: Minimum Viable Stack

Before the economic layers (4–5), this section consolidates the **minimum viable requirements**:

Layer	Minimum Viable Requirement	Falsification Test
L0	At least one L0-B grade profile; sampling for L0-A	If proving concentrates on L0-A hardware and lot-attestation coverage stays below the L0-B threshold ($\geq$80% of production lots) → thesis weakened
L1	Encrypted P2P + $\geq$2 fallback transports (at least one look-like-something class); VerifyReach with $\geq$3 operators	If >20% of sampled ASNs *in uncensored regions* (outside the observatory's pre-declared censored-region list, Section 18.5) are unreachable for >7 days → thesis fails
L2	Threshold-signed releases + transparency log + $\geq$2 builders	If single key compromise can push malicious updates → thesis fails
L3	$\geq$3 identity predicates with $p_{95,t}$ of the canonical tuple <5s; issuer pluralism	If identity requires centralized issuer → thesis weakened

Table 20.3: Minimum viable requirements for Layers 0–3.

Telemetry Reproducibility Principle

Telemetry must be reproducible from raw public artifacts by multiple independent parties, or it is not part of the trust base.

Specifically:

1. **Raw data availability:** All inputs to telemetry metrics are published or archived.
2. **Open methodology:** Aggregation formulas and statistical procedures are documented.

3. **Multi-party verification:** At least 3 independent parties publish telemetry; divergence >10 percentage points (absolute, between independently computed aggregates) triggers investigation.

Four layers now stand underneath the economic ones: matter that can be audited, packets that survive pressure, code that ships without a gatekeeper, and identity that travels without a name. Each was specified so that its failure is a measurement, not a debate. Part IV asks the only question left for a stack this deep: what work is worth settling on it.

Part IV
Truth, Work, and Settlement: Layers 4–5

Part III started at the basement: Layer 0 as verifiable machines and energy. Layers 1–3 keep packets moving, code running, and actors named without being doxxed.

Part IV is where the stack becomes **explicitly economic**. Layers 4 and 5 are the point where:

- raw compute becomes **truth** (via proofs and verification asymmetry), and
- truth becomes **settlement** (via private, non-custodial value flows).

This is the layer pair that turns the abstract triad (Privacy, Proofs, Compute) into something you can actually buy and sell: **proof commodities, verified FLOPs, and privacy-preserving flows**.

Chapter 21

From Infrastructure to Economics

Layers 0–3 answer four questions:

1. *Can we trust the machine?* (Layer 0)
2. *Can we reach it under censorship?* (Layer 1)
3. *Can we get code onto it and keep it updated?* (Layer 2)
4. *Can actors prove "who" they are without doxxing themselves?* (Layer 3)

Layers 4 and 5 take those answers as given and ask two more:

5. *Can we turn machine work into claims that anyone can verify cheaply?* (Layer 4)
6. *Can we settle those claims as value flows without chokepoints?* (Layer 5)

If Layers 0–3 are the **nervous system and limbs**, Layers 4–5 are the **cortex and circulatory system**: they decide what counts as action and how it is remembered in economic form.

Infrastructure $\neq$ Money

Layers 4–5 create markets for proofs, verified compute, and private settlement. They do not create money at all. A working market for verified FLOPs is necessary for the monetary thesis but not sufficient: Part II's value-capture conditions decide whether the market's economics *accrue* to the base asset through fees, burns, and collateral (Value Capture Lemma), and whether accrual becomes moneyness is the separate question of the three link-9 mechanisms.

Chapter 22

Layer 4: Truth & Work

Layer 4 is the **Prove/Verify spine** of the stack.

Its job is to take arbitrary claims—"this media object came from camera C at time T," "this inference was computed by model M on input X," "this corridor swap executed atomically"—and turn them into **portable proofs** that:

- anyone can check with **cheap, public verification**, and
- can be **standardized as commodities** (canonical workloads, SLAs, prices).

22.1 What Layer 4 Is (and Is not)

Layer 4 is:

- The layer of **circuits, proofs, and verification economics**.
- The place where **verification asymmetry** is engineered and measured.
- The home of **Proofs-as-a-Library (PAL)**, proof factories, and **canonical workload registries**.

Layer 4 is not:

- A specific proof system (SNARK vs. STARK vs. something else); it assumes **multi-ZK**.
- A single PoUW design; it supports several patterns as long as VERIFYPRICE and anti-capture constraints are met.
- A truth oracle. Proofs do not prove truth; they prove specific claims under stated assumptions—origin, custody, computation, policy compliance, and settlement finality. The thesis is not that cryptography tells us what is true, but that it reduces the surface area over which institutions must be trusted.

22.2 Verification Asymmetry Revisited

Recall the definition from Part I. For a workload W:

- $p(W)$ = cost (time, energy, hardware) to **produce** a result + proof.
- $v(W)$ = cost to **verify** that result + proof.
- $r(W) = \frac{v(W)}{p(W)}$ = **verification asymmetry**.

Both $v(W)$ and $p(W)$ are aggregated into the same numeraire — USD via the per-run price

model of Section 22.6 ($c_{\text{run}} = (e_{\text{run}}/3.6 \times 10^6)\, p_{\text{energy}} + t_{\text{run}}\, p_{\text{opp}}$) — so the ratio is dimensionless and both halves carry the same measurement contract: v under the adversarial harness, p under the reference production model (Appendix A). Read the denominator before celebrating the ratio: $r(W)$ improves in two ways, verification getting cheaper or production getting dearer, and a passing $r(W)$ achieved by expensive production is a cost problem wearing an asymmetry credential. The ratio is therefore published with both components; a falling $r(W)$ driven by a rising $p(W)$ is flagged as such.

Layer 4's goal is simple:

> **Make $r(W) \ll 1$ for the workloads that matter, and keep it that way in production.**

This is the touchstone of the gold analogy, now carrying a price: expensive to smelt, cheap to assay. $r(W)$ is the assay, priced.

Why this matters:

- If $v(W)$ is small and stable, **anyone** (including small nodes) can check proofs.
- That makes proofs and verified FLOPs **commodities**: units of work that any counterparty can accept without trusting a platform.

Layer 4 introduces:

$$\text{VerifyPrice}(W) = \{\underbrace{\text{p50 time},\ \text{p95 time},\ \text{p50 cost},\ \text{p95 cost},\ \text{failure rate}}_{\text{canonical five-field tuple } (\textit{Appendix A})},\ \text{hardware profile mix}\}$$

where the first five fields are the canonical tuple of Appendix A and the hardware profile mix is Layer 4's reporting context: every published series names the verifier classes it aggregates over, and the mix travels with the tuple rather than replacing any field.

22.2.1 Verification Modalities, and Which One Carries the Monetary Claim

The ratio $r(W)$ has been carried through this document as a single scalar against a single ceiling. That was adequate while cheap verification was an operational requirement. It is not adequate now that Section 13.3.10 makes cheap verification the mechanism the monetary claim runs on, because **the ceiling as written does not protect the property the monetary claim needs.**

The ceiling is not binding where it matters. A workload can satisfy $r(W) \leq 0.3$ and be nowhere near the regime the touchstone image describes. At $r = 0.3$ checking a claim costs thirty per cent of producing it. Nobody checks everything at that price. Participants check samples, or delegate, or accept a counterparty's assurance — which is precisely the *symmetric ignorance* described in Section 13.3.10, arrived at by an SLO the network is passing. An asset whose verification sits at the ceiling is accepted because checking is too expensive to bother

with, which is the old mechanism with better tooling, and it fails the way the old mechanism fails.

The ceiling is therefore an engineering floor and should be labelled as one. What the monetary argument requires is a statement about *modality*, because the three ways of establishing that work was done correctly differ by orders of magnitude and differ in kind.

Verification Modality Bands

M1 — Succinct. A cryptographic proof checked by a verifier whose cost is sublinear in, and typically independent of, the work proved. Target $r(W) \leq 10^{-3}$, stretch 10^{-5}. *Everyone can afford to check everything.*

M2 — Algebraic or probabilistic. A structured check cheaper than re-execution by a factor of the problem size, accepting a stated soundness error — Freivalds-style matrix verification [Freivalds 1977] at $O(n^2)$ against $O(n^3)$ production is the canonical case. Each Freivalds round costs three matrix–vector products, so k rounds at soundness 2^{-k} give $r(W) \approx 3k/n$ in arithmetic alone (a single round over a large field, with soundness $1/|\mathbb{F}|$, gives $\approx 3/n$); the $O(n^2)$ input/output and hashing cost of reading the committed operands adds a constant. Target $r(W) \leq 10^{-2}$ with the soundness parameter published beside the ratio. *Everyone can afford to check everything, in expectation.*

M3 — Replicated or attested. Correctness established by independent re-execution, committee agreement, or hardware attestation. Verification cost does not fall asymptotically below production cost; what is bought is fault tolerance and accountability, not asymmetry. No $r(W)$ target applies, because the ratio is not the operative quantity. *Someone checked. You are trusting who.*

The aggregate $r(W) \leq 0.3$ ceiling is retained as a portfolio-level engineering floor and as the fallback for workloads not yet assigned a band. It is not a monetary threshold and no monetary claim reads on it.

Only M1 and M2 carry the monetary mechanism, and this is a consequence rather than a preference. Section 13.3.10 locates moneyness in no-questions-asked acceptance reached by symmetric knowledge — checking is free enough, and public enough, that no counterparty can acquire an informational edge by performing it. M1 delivers that, to a security level published in provable rather than conjectured bits. M2 delivers it up to a published soundness error, which is a quantified and therefore acceptable dilution. M3 does not deliver it at all: under replication the participant is not checking, a committee is, and the participant is trusting the committee's composition, independence, and continued availability.

That is not a criticism of M3. Replication is the right construction for a great many workloads and much of the useful capacity this thesis describes will be M3. **It is a statement about what M3 capacity can be used for.** M3 supports service delivery, Work Credits, and SLA tiers.

It does not support the pledgeability claim, because a position whose correctness rests on a committee requires the lender to form a view about that committee — which is due diligence, which is the thing information-insensitivity is defined as the absence of.

> Replication buys fault tolerance. Succinctness buys moneyness. The thesis needs both and has been pricing them as though they were the same good.

The modality is what the haircut was always pricing. This resolves a loose end rather than creating one. The collateral haircut schedule of Section 17.4 already discounts positions by how legible their substrate is, and Section 13.3.10 identifies those haircuts as prices on residual unit-quality uncertainty. Modality is the second axis of the same schedule and the more fundamental one: an M3-verified position carries irreducible delegated-trust exposure regardless of how good its hardware grade is, and an M1-verified position on modest hardware may carry less. The haircut table is therefore indexed on *modality* $\times$ *grade*, and the two are not collapsed: Table 17.3 in Section 17.4.2 is the joint table, and the modality increment exceeds the grade increment in every cell but one.

What this changes in the instruments. Three reporting consequences, none of which relaxes anything:

- VERIFYPRICE series are published **banded by modality**. An aggregate $r(W)$ across a mixed M1/M3 portfolio is a number with no interpretation, and its improvement can be produced by changing the mix rather than by improving anything.
- The **M1+M2 share of verified capacity** becomes a published series in its own right. It is the fraction of the network's output that carries the monetary mechanism at all, and it belongs beside fee coverage on the Value Capture Board (Section 26.9) rather than among the engineering metrics.
- Red Line 3's monoculture condition acquires a second reading. Concentration of verification in a single vendor is one failure; **migration of the capacity mix from M1 toward M3** is a quieter version of the same failure, because both end at "trust the verifier." The second reading is instrumented by Invariant 1(a) of Section 25.4: the M1+M2 share of verified units is a published series with a floor, and a fall below it is Sev-1.

And the rhetoric can now be correct. Claims that checking costs a small fraction of one per cent of producing are true of M1 and of M2 at realistic problem sizes — one Freivalds round at $n = 4096$ costs roughly seven parts in ten thousand of production ($3/4096$), and twenty rounds at soundness 2^{-20} cost about 1.5×10^{-2}, which is *outside* the M2 band. The band therefore admits `MATMUL_4096` under Freivalds only at $k \leq 13$ rounds ($\approx 9.5 \times 10^{-3}$, soundness 2^{-13}) or under the large-field single-round variant, and the registry entry must say which; larger n relaxes the constraint linearly. An earlier draft quoted "two parts in ten thousand"

here, which is $1/n$ and counts one matrix–vector product where the check needs three. The claims are false of M3 and of the aggregate ceiling. Any statement of the form "checking costs one ten-thousandth of producing" is a claim about succinct verification specifically and must name the modality wherever it appears, in this document and in every derivative artifact. The looser figure was never wrong as a floor. It was wrong as a description, and it was quoted as a description.

22.3 Canonical Workloads and Proof Types

To avoid an unbounded zoo of bespoke proofs, Layer 4 maintains **canonical workloads**:

- **MatMul(**$n, k, m; \varepsilon_r$**)** – matrix multiplication at specified dimensions and componentwise relative error bound.
- **Inference(**M, X**; policy)** – run model M on input X under constraints.
- **Provenance(**C**, chain)** – provenance chain for content type C.
- **Settlement(**S**, policy)** – settlement and refund logic for a corridor.

For each canonical workload, the stack defines:

- **Circuits / arithmetization**: how the workload is represented for proving systems.
- **Proof schemas**: which proof systems are supported.
- **Verification modality**: M1, M2, or M3 (Section 22.2.1); a computation that can be checked under two modalities is registered as two SKUs, never as one SKU carrying both.
- **SLO tiers**: latency and assurance classes (Section 22.3.2).

Canonical workloads matter economically:

- They become the **SKUs** for proof and compute markets.
- Work Credits are sold against committed capacity for units of these workloads.
- VERIFYPRICE and $r(W)$ are tracked per workload and tier.

Every SKU this document names lives in one table, the workload registry of Section 22.3.3; the definition template, the starter set (Section 22.10), the constitutional SLO table (Section 22.6), and the primitive catalog (Section 24.1) all point at it rather than carrying their own lists.

The precedent for standardized work. The requirement that demand standardize into canonical units before it can become an economic object has a long pedigree, and it is worth naming because it is easy to mistake for an engineering preference. The standardization thesis in the print literature [Eisenstein 1979] documents the sequence: typographic uniformity—every letter identical to every other letter of its kind, modular, movable, recombinable—was daily practice at scale for roughly three centuries before the machine shop embodied the same template in screws, gauges, and interchangeable parts. The claim is causal and well-documented as mechanism: standardization of the cognitive module preceded and enabled standardization of the physical one. The stronger claim—that industrialization was possible *only* because of

movable type—is not adopted here. The analogy to this layer is exact at the level of function: canonical workloads and PIDL receipts are what make two providers' units of work identical, comparable, and combinable; diffuse demand for "compute" becomes an economic unit only when the units are interchangeable. The registry is the type case; the receipt is the printed page.

22.3.1 Canonical Workload Definition Template

For a workload to become a tradable SKU (and eligible for Work Credit issuance), it must be defined precisely enough that two providers' units are interchangeable—and that is the whole reason this template exists. Free-form compute cannot be a unit of account. If one vendor's "matmul" means fixed dimensions under a relative error bound and another's means whatever their kernel happens to produce, no counterparty can compare the two, no market can price them, and no registry can mint against them. The template fixes the fields that turn work into a unit: what is claimed, what is public, what is private, what gets checked, what it costs to check, and where it may run.

One field deserves to be stated in the open rather than left inside a parenthesis, because two natural drafts of this template get it wrong in opposite directions. The error bound is *relative, componentwise, and normalized to the operands that produced each entry* — never absolute, and never a matrix-norm inequality. An absolute bound fails honest provers: standard BLAS kernels reorder floating-point arithmetic, and the discrepancy that reordering introduces grows with the magnitude of the operands, so a correct product computed by a different honest kernel would still fail the test. A matrix-norm bound fails in the other direction. Written as $\|C - A \cdot B\|_\infty \le \varepsilon_r \|A\|_\infty \|B\|_\infty$ with the matrix ∞-norm (maximum row sum), the right-hand side is a per-row budget of order $\varepsilon_r \cdot n^2 \cdot \mathbb{E}|a|\,\mathbb{E}|b|$. At $n = 4096$ on i.i.d. Gaussian FP32 inputs with $\varepsilon_r = 10^{-3}$ that budget is about 1.16×10^4 against an honest FP32 error of 0.23: a product computed from INT8-quantized inputs errs by about 3.97×10^3 and passes, and a single output entry corrupted by about 1.05×10^4 (typical $|C_{ij}| \approx 64$) also passes. A bound that admits the next-cheaper precision class and a two-hundred-fold corruption of one entry is not an acceptance test; it is a formality. The correct form is the componentwise backward-error bound of floating-point analysis [Higham 2002] (Higham, *Accuracy and Stability of Numerical Algorithms*):

$$|C_{ij} - (A \cdot B)_{ij}| \;\le\; \varepsilon_r \sum_k |A_{ik}|\,|B_{kj}| \quad \text{for every } (i, j),$$

with ε_r set from the pinned reference kernel's precision — $\varepsilon_r \approx 10^{-5}$ for strict FP32 accumulation — and calibrated so that the next-cheaper precision class the template forbids (TF32, FP16, INT8) fails. The relative componentwise bound forgives honest implementations and still fails results that are actually wrong, entry by entry. Strict with dishonesty, tolerant of implementation—that is the register the entire template is written in, and a bound has to be

tested against dishonesty before it earns the register: every workload registration ships an adversarial acceptance suite (INT8 input substitution, single-entry corruption, row zeroing, precision downgrade of the accumulator) that the bound must reject, and a registration whose bound passes any of them is returned.

Canonical Workload Definition Template

WorkloadID: Unique identifier from the registry of Section 22.3.3 (e.g., `MATMUL_4096_FP32_FREIVALDS`, `MATMUL_4096_FP32_SNARK`, `INFER_LM_70B_256TOK`). One ID names one modality; the same computation under a second modality is a second ID.

Statement: What is being proved (natural language + formal predicate)

- *Example:* "The matrix product $C = A \times B$ was computed correctly, where $A, B \in \mathbb{R}^{4096\times4096}$, and $|C_{ij} - (A\cdot B)_{ij}| \le \varepsilon_r \sum_k |A_{ik}|\,|B_{kj}|$ for every (i,j)" — the *componentwise, scale-normalized* bound of Section 22.3. Neither an absolute bound nor a matrix-norm bound is the template's form: the first fails honest BLAS kernels, the second admits INT8 execution.

Public Inputs:

- Commitment to A, B (hash or Merkle root)
- Claimed result commitment (hash of C)
- Error bound ε_r and the reference kernel ID that sets it
- Timestamp range

Private Inputs (Witness):

- Full matrices A, B, C
- Intermediate computation trace (if required by proof system)

What Is Verified:

- Correctness (computation matches claim)
- Bounds (result within specified limits)
- Freshness (timestamp within allowed window)
- Liveness (proof generated within epoch, not precomputed)

Adversarial Acceptance Suite: the named set of dishonest transcripts the acceptance test must reject (for MatMul: INT8 input substitution, TF32/FP16 accumulation, single-entry corruption, row zeroing). Registration is refused while any suite member passes.

Verifier Complexity Class:

- Time: two regimes, and the registry entry must say which it prices. *Succinct-proof* (M1) workloads verify in $O(\log n)$ after the committed output is in hand; *Freivalds-style* (M2) audit workloads verify in $O(n^2)$ — reading and hashing the committed operands, then k matrix–vector rounds — but avoid the $O(n^3)$ recompute. The template's complexity field names the regime. A computation checkable in both regimes is two SKUs with two entries (`MATMUL_4096_FP32_FREIVALDS` and `MATMUL_4096_FP32_SNARK` in Sec-

tion 22.3.3); an entry never publishes both regimes under one ID, because the two have different modalities, tiers, and prices.
- Memory: specified peak (e.g., ≤512MB for Laptop-Class)
- Proof size: max allowed (e.g., ≤1MB)

Policy Hooks: (machine-verifiable predicates)
- Hardware profile requirements: the minimum grade in the Allowed Hardware Profiles field below (L0-A by default), with a higher floor named per SLO tier where the entry requires one (e.g., L0-B or higher for the Gold latency tier)
- Prover stake requirements
- Membership/non-membership predicates (allowlist proofs, not graph inspection)

Allowed Hardware Profiles:
- L0-A: Yes (with issuance weight 0.9×)
- L0-B: Yes (1.0×)
- L0-C/D: Yes (1.0×; open-hardware premium via collateral grade and routing priority, never via issuance above delivered work)

22.3.2 SLA Tiers: Latency and Assurance

The obvious tier table — Bronze/Silver/Gold at 60 s/10 s/2 s prove-side latency with 1×/2×/3× verification — prices the wrong thing. "3× verification + audit" is replication: modality M3 in the taxonomy of Section 22.2.1, the one modality that carries no monetary mechanism. A 2 s prove-side target is reachable today only under M2 or M3, because succinct (M1) proving of any workload in the registry runs minutes to hours on current provers. Its highest-fee tier would therefore rest on the least monetary modality and reward provers for choosing it. Latency and assurance are separate goods with separate prices, and the tiers are two orthogonal axes: a latency tier that prices how fast the attested result arrives, and an assurance tier that prices how it was verified. A quote names one of each.

Latency tier	Prove-side p95 to attested result	Fee multiplier
Bronze	≤ 60 s	1.0×
Silver	≤ 10 s	1.3×
Gold	≤ 2 s	1.6×

Table 22.1: Latency tiers for canonical workloads. The target is the *prove-side* time to produce the attested result (generation plus attestation or proof), not verify-side time: verification SLOs are constitutional and set separately in Section 22.7, and no prove-side tier can relax them. Gold is attainable today only with M2 or M3 assurance; an M1 quote at Gold latency is refused by the registry until a prover benchmark on record supports it.

Assurance tier	Modality	What the buyer receives	Fee multiplier
Replicated	M3	One independent re-execution or attestation; committee is the trust root	1.0×
Replicated + audit	M3	Three independent re-executions plus randomized third-party audit; Tier C/B service claims	1.4×
Probabilistic	M2	Algebraic check with published soundness parameter (e.g., Freivalds $k = 13$, error 2^{-13}); Tier B Work Credits	1.6×
Succinct	M1	Cryptographic proof checked in milliseconds independent of workload size; Tier A Work Credits when the maturity gate lifts. Prove-side latency SLO set per workload from the current prover benchmark on record, not from the latency table	2.5×

Table 22.2: Assurance tiers for canonical workloads, indexed by verification modality (Section 22.2.1). The highest multiplier attaches to the modality that carries the monetary mechanism, and its latency promise is whatever current provers can honestly meet. A quote's total multiplier is the product of its latency and assurance multipliers.

Example: MATMUL_4096_FP32, two SKUs The same computation is registered twice, because it is checkable under two modalities and the two are different goods.

Why this matters: Without this template, "canonical workload" risks becoming "whatever the prover says it is." With it, workloads are **precisely specified** (anyone can implement a conforming prover/verifier), **auditable** (claims can be checked against the template), and **tradable** (markets can price and exchange standardized units).

22.3.3 Workload Registry

A registry that is quoted three ways is not a registry. Table 22.4 is the single table; every other mention of a SKU in this document is a pointer into it, and an identifier that does not appear here is not a canonical workload.

22.4 PoUW Design Patterns

Layer 4 does not pick a single consensus recipe. It supports several **proof-of-useful-work patterns**, as long as they satisfy:

- **Open admission**: commodity participants can join.
- **Unpredictable leader election**: no one can cheaply bias the lottery.
- **Useful work binding**: you cannot precompute offline or reuse stale work.
- **Proof quality & anti-spam**: junk proofs cannot flood the system without penalty.

Pattern 1: Hash-gated useful work

Field	`MATMUL_4096_FP32_FREIVALDS`	`MATMUL_4096_FP32_SNARK`
Statement	Matrix multiply $C = A \times B$, dimensions 4096×4096, FP32; componentwise bound $\lvert C_{ij} - (A \cdot B)_{ij} \rvert \le \varepsilon_r \sum_k \lvert A_{ik} \rvert \lvert B_{kj} \rvert$ for every (i, j), $\varepsilon_r = 10^{-5}$ (strict FP32 accumulation, pinned kernel)	
Modality / tier ceiling	M2 / Tier B	M1 / Tier A (maturity-gated; Section 22.10)
Public Inputs	Hash(A), Hash(B), Hash(C), ε_r, reference kernel ID, timestamp	
Private Inputs	A, B, C	A, B, C, arithmetization trace
Verified	Correctness (probabilistic, soundness 2^{-k} published), Bounds, Freshness	Correctness, Bounds, Freshness
Verifier work	$O(n^2)$: stream and hash A, B, C against the public commitments (192 MB of FP32 operands), then k rounds of three matrix–vector products with the componentwise tolerance applied per row; $k = 13$ default ($r \approx 9.5 \times 10^{-3}$)	$O(\log n)$ proof check after the committed output is in hand; the verifier never touches A or B. The buyer still receives C (∼64 MB) and hashes it once against the commitment
Memory	≤256MB with streaming (operands read once per round, never held whole); a verifier that materializes all three matrices needs ≥192MB and sits at the ceiling	≤256MB
Proof / receipt size	Receipt only (≤500KB); no proof artifact	≤500KB proof plus receipt
Prove-side feasibility	Production: the prover computes C once with the pinned kernel	Research: no production prover exists for a 4096^3 FP32 product at a price a buyer would pay; registered so the target is public
L0 Requirement	L0-A minimum; L0-B+ for Gold latency	L0-A minimum
VerifyPrice Target	$t_{95} \le 5$s (Laptop-Class; I/O-bound on the 192 MB read)	$t_{95} \le 2$s (Laptop-Class)

Table 22.3: Example canonical workload: the two MATMUL_4096_FP32 SKUs. On the tolerance: FP32 accumulation over $k = 4096$ terms has typical rounding error of order $\sqrt{k}\,u \approx 4 \times 10^{-6}$ and worst case $k\,u \approx 2.4 \times 10^{-4}$ relative to $\sum_k \lvert A_{ik} \rvert \lvert B_{kj} \rvert$ (unit roundoff $u = 2^{-24}$); $\varepsilon_r = 10^{-5}$ therefore accepts every honest strict-FP32 kernel in the typical case and is calibrated to reject TF32 accumulation, whose 10-bit mantissa puts its error near 10^{-3}. A tolerance of 10^{-4}–10^{-3} is characteristic of TF32, which the template forbids, not of FP32. The reference kernel ID pins the arithmetic order so heterogeneous provers produce comparable transcripts; see the determinism note in Section 22.3.

WorkloadID	Class	Modality	Tier ceiling	Acceptance criterion	Reference verifier class
PROOF_2^20	Proof	M1	A	Succinct proof of a 2^{20}-constraint statement verifies against the published verifying key	Laptop
MATMUL_4096_FP32_FREIVALDS	MatMul	M2	B	$k = 13$ Freivalds rounds pass the componentwise bound $\varepsilon_r = 10^{-5}$ per row; operands hash to the public commitments; soundness 2^{-13} published	Laptop
MATMUL_4096_FP32_SNARK	MatMul	M1	A (gated)	Succinct proof that C satisfies the componentwise bound against committed A, B; buyer hashes delivered C	Laptop
INFER_LM_7B_512TOK	Inference	M3 (sampled re-execution)	B	Pinned-kernel transcript; $\geq 10\%$ of inferences re-executed by random auditors; divergence beyond ε_r fails; stake slashed on fraud	Server (Laptop re-execution is minutes; published, not constitutional)
INFER_LM_7B_512TOK_ZKML	Inference	M1	A (gated)	Succinct proof of the quantized forward pass under committed weights that hash to the published model; quantization gap to the advertised model published separately	Laptop
INFER_LM_70B_256TOK	Inference	M3	B	As INFER_LM_7B_512TOK; re-execution bound to Server-Class (model does not fit Laptop-Class memory)	Server
PROVENANCE_MEDIA_V1	Provenance	M1	A	Proof that the asset's hash chain is rooted in a Layer-0 attested capture and each hop is signed by the stated party under an unrevoked certificate	Laptop, Mobile
SETTLE_ATOMIC_BTC_XMR_V1	Settlement	M1 receipt over chain finality	A	Both legs finalized, or the no-loss exit of Section 23.4 taken within the refund window; receipt carries both chain proofs	Laptop
SETTLE_ATOMIC_BTC_ZEC_V1	Settlement	M1 receipt over chain finality	A (gated)	As BTC↔XMR, over the shielded ZEC pool; *registered, not admissible*: no production shielded-ZEC adaptor-signature swap exists (Section 23.3)	Laptop

Table 22.4: The workload registry. "Tier ceiling" is the highest Work Credit tier the SKU can reach once its verification component is Production-mature (Section 22.10); "gated" marks SKUs whose prove side is research-grade today and which therefore issue at Tier B or below, or not at all, until the maturity table lifts. Modality per Section 22.2.1. The reference verifier class is the one the constitutional SLO is measured on (Section 22.6); where it is Server-Class, the Laptop-Class figure is still published.

- Miners perform a cheap hash race; crossing a threshold gives short-lived eligibility.
- To produce a valid block, the miner attaches a PoUW artifact seeded from header randomness.

Pattern 2: Proof-first selection

- Provers race to produce useful-work proofs and post them to a mempool.
- A lightweight mechanism selects which proofs get included and rewarded.

Both patterns are compatible with MatMul-PoUW, Inference-PoUW, and hybrid schemes.

22.4.1 The Usefulness–Unpredictability Dilemma

Neither pattern resolves the problem that sits under every proof-of-useful-work design, and this document has until now stated the problem's two halves in separate places as though each were solved. It is one problem, it is open, and it should be named as such.

A work function must be *unpredictable* to secure consensus: if the input is known in advance, a miner can precompute and the lottery is not a lottery (property P1 below). A work function must be *useful* to have a buyer: someone must want the result of *this* input, not a random one. The two requirements pull on the same field of the block. Header-seeded inputs — Pattern 1 as written — are unpredictable and useless: nobody buys the product of two matrices drawn from a block hash, so the "useful" work is heat with better branding, and the demand curve the thesis relies on (Section 22.10) does not exist for it. Client-supplied inputs — the only kind a buyer pays for — are predictable to whoever supplied them: a client and a miner who are the same party, or who collude, can precompute the work, submit it as fresh, and bias the lottery by exactly the amount P1 forbids. The hybrid designs — random-seeded verification challenges layered on client work, in the manner of Ambient's proof-of-logits [Ambient 2024] — move the randomness from the input to the *audit*: the client's input is taken as given and the header seeds which slices of the work get checked. That is the right shape, and it is not a solution until it comes with a stated collusion bound: for a client–miner coalition controlling fraction ϕ of demand and fraction θ of hashrate, what reward advantage does it obtain, and under what audit rate does that advantage fall below the grinding tolerance δ of P2? No design this document cites publishes that bound.

This is the flagship research problem for a build lab, and it is stated here as one rather than folded into an engineering table. Verification asymmetry is a theorem [Freivalds 1977] and is not the hard part. The hard part is making the work simultaneously worth buying and worth mining, and the honest current status is: unsolved, with a promising shape and no proof. Until a hybrid design carries a collusion bound, MatMul-PoUW and Inference-PoUW secure the *service market* (Work Credits sold against reserved capacity and extinguished by delivered, verified work) and do not by themselves secure *consensus*; the consensus role remains a conjecture in Table 22.5, and the starter-set rows inherit that label.

PoUW Security Properties (Must Hold)

P1. Precomputation Resistance: Proofs seeded by unpredictable randomness; valid only for $N \leq 10$ blocks after seed.

P2. Grinding Resistance: Attacker with $X\%$ hashrate gains at most $(1+\delta) \cdot X\%$ rewards, where $\delta \leq 0.1$.

P3. Network Advantage Bound: p95 propagation delay $\leq$ 2s for proof announcements.

P4. Spam Resistance: Deposit $\geq 10\times$ expected verification cost; 100% slash for invalid proofs.

P5. Cartel Detection: Top-1 share $< 20\%$; Top-5 share $< 50\%$; entry latency $\leq$ 7 days.

P6. Useful Work Binding: $\geq 10\%$ of proofs independently re-verified by random auditors.

Component	Maturity	Status
MatMul verification asymmetry (Freivalds, $O(n^2)$ check)	Production	Theorem [Freivalds 1977] (Freivalds 1977); the asymmetry itself is not in question
MatMul-PoUW consensus security (grinding / precomputation bounds under useful inputs)	Experimental	Conjecture; no published collusion bound (Section 22.4.1); deployed on a public mainnet since 2026 [Basu 2026]
Inference PoUW (Proof-of-Logits)	Pilot	Tested in limited networks
Succinct MatMul proving at 4096^3 FP32 (`_SNARK` SKU)	Research	No production prover at buyer-acceptable cost
Full ZKML for large models	Research	Not yet practical at production scale
Hash-gated useful work (Pattern 1)	Experimental	Devnet demonstrations
Proof-first selection (Pattern 2)	Experimental	Early implementations

Table 22.5: PoUW component research maturity.

Research Maturity Labels: **Implication for service claims:** The "AI Money" label is an analytical lens on verified-compute demand, not an instrument classification. As components mature from Experimental to Production, inference Work Credits can receive stronger assurance grades and smaller service haircuts. Their issuance remains separate from the base-asset monetary constitution (see Tier taxonomy in Section 24.1).

22.5 Proof Factories and PaL

Most developers do not want to think about circuits; they want to say:

> *"Prove that this computation happened, then get me a receipt and pay whoever did the proving."*

Layer 4 provides this via:

Verification Type	Service-Claim Eligibility
Full cryptographic proof (Tier A)	Full eligibility for high-assurance Work Credit issuance
Probabilistic audit with explicit error bounds (Tier B)	Discounted or limited eligibility
TEE/vendor attestation only (Tier C)	Service market only; not eligible for high-assurance capacity claims
Unverified output	No eligibility

Table 22.6: Verification type and typed service-claim eligibility.

- **Proof factories**: infrastructure clusters specialized in generating proofs.
- **Proofs-as-a-Library (PAL)**: an SDK that compiles high-level claims into proofs.

  PAL exposes interfaces like:

```
prove_compute(f, inputs, policy)
prove_provenance(asset_id, lineage)
prove_settlement(tx, corridor_policy)
```

Under the hood, PAL:

1. Maps the request to a canonical workload W.
2. Selects suitable proof systems and hardware profiles.
3. Submits the job via neutral routers to proof factories.
4. Returns a PIDL receipt plus a proof artifact.

22.6 VerifyPrice Observatory

Verification asymmetry is a design goal; VERIFYPRICE turns it into a dashboard.

The **VerifyPrice observatory** continuously measures:

- p50/p95 verify times,
- estimated energy per verification,
- failure rates and mismatch rates,
- diversity metrics (how many independent verifiers are actually checking).

Class	Hardware Spec	Use Case
Laptop-Class	4-core CPU, 16GB RAM, SSD, no GPU	Baseline for "anyone can verify"
Mobile-Class	ARM SoC, 8GB RAM, flash	Edge verification, IoT
Server-Class	32-core CPU, 128GB RAM, optional GPU	High-throughput

Table 22.7: Reference verifier classes for VerifyPrice measurement.

Reference Verifier Classes:

22.6.1 VerifyPrice Measurement Specification

VerifyPrice is a service-assurance KPI and one necessary input to the base asset's conditional monetary test. It determines whether Work Credits retain a verifiable service proposition; it does not make them monetary. That requires a rigorous measurement harness, not just a dashboard slogan.

The harness exists because the observatory is not a disinterested instrument. It is a publisher with interests, run by operators who can be pressured, and the numbers it publishes are exactly the kind a protocol under strain would prefer not to publish. Assume the observatory is fallible, and further assume it is *adversarially* fallible—inclined toward flattering measurements when the truth is inconvenient. The specification below is not written to make the observatory honest. It is written so that when the observatory lies, anyone can catch it: the samplers are named, the harness is open, the raw data is archived, and independent operators run the same tests on the same proofs. A lie has to survive all of that. The spec does not produce trust in the measurement; it produces detection of the absence of it.

Input View (What the Harness Records): VERIFYPRICE has exactly one definition: the canonical five-field tuple of Appendix A, introduced at Section 22.2,

$$\text{VerifyPrice}(W) = (p_{50,t},\ p_{95,t},\ p_{50,c},\ p_{95,c},\ \text{fail}).$$

Tuples that resemble it, here and in Section 26.3, are the harness's inputs, not rival definitions wearing the output's name. What the harness records, per run, for workload W on verifier class V, is the **input view**:

$$\text{Inputs}(W, V) \ni (t_{\text{run}},\ e_{\text{run}},\ m_{\text{run}},\ b_{\text{run}},\ \text{outcome}_{\text{run}})$$

Where:

- t_{run}: verification **time** (seconds); $p_{50,t}$ and $p_{95,t}$ are its percentiles over the window
- e_{run}: **energy** per verification (Joules, measured via hardware counters or watt-meter); enters c_{run} below
- m_{run}: peak **memory** (MB) — *published, not priced*: a run exceeding the registry entry's memory ceiling is recorded as a failure and lands in fail; observed peak is published beside the tuple
- b_{run}: **bandwidth** consumed (bytes) — *published, not priced*, on the same rule as memory
- $\text{outcome}_{\text{run}}$: pass, or a failure code (timeout, invalid proof rejected, crash, resource ceiling exceeded); fail is the failure fraction

USD cost is derived *per run*, then percentiled: $c_{\text{run}} = \dfrac{e_{\text{run}}}{3.6 \times 10^6} \times p_{\text{energy}} + t_{\text{run}} \times p_{\text{opp}}$, where e_{run} is the run's measured energy in Joules, p_{energy} is the reference energy price (USD 0.10/kWh, so the first term divides by 3.6×10^6 J/kWh), and p_{opp} is opportunity cost (USD

0.001/s), both published quarterly. The cost percentiles of the canonical tuple — $p_{50,c}$ and $p_{95,c}$ — are the percentiles of c_{run} across the measurement window, never recomputed from percentile inputs; in particular, deriving a cost figure from t_{95} alone is not the construction.

Requirement	Rationale
$\geq$2 independent implementations per workload tier	Prevents single-implementation bugs from corrupting measurements
Open-source, reproducible builds	Anyone can audit and rebuild
Deterministic output	Same proof $\rightarrow$ same result, always
Versioned and tagged	Measurements tied to specific verifier version

Table 22.8: Verifier implementation requirements.

Verifier Implementations:

Sampling Methodology:

- **Random selection:** Proofs selected uniformly at random from recent submissions (not cherry-picked).
- **Stratified by region/profile:** Measurements cover $\geq$10 regions and $\geq$3 hardware profiles per workload.
- **Adversarial corpus:** 10% of test proofs are malformed or worst-case (max witness size, pathological inputs) to measure failure handling.
- **Continuous measurement:** Not periodic snapshots; rolling 24-hour windows published hourly.

Condition	How Tested
Network latency	200ms RTT injected for proof fetch
Packet loss	5% random packet loss during verification
Worst-case proof size	Max allowed witness for workload class
Malformed proofs	10% of test corpus intentionally invalid
Resource exhaustion	Verification under 80% memory pressure

Table 22.9: Adversarial conditions for VerifyPrice testing.

Adversarial Conditions:

Reproducibility Standard:

- **Benchmark harness:** Open-source, versioned, deterministic. Anyone can run the same tests.
- **Signed results:** Each measurement batch is signed by $\geq$2 independent measurement operators.

- **Divergence alerts:** If independent operators diverge by >5%, investigation triggered.
- **Archived raw data:** All proof samples and timing logs archived for 1 year.

Why this matters: Without this spec, "VerifyPrice observatory" is an oracle claim. With it, VerifyPrice becomes **reproducible consensus**—any skeptic can run the harness, check the measurements, and falsify the dashboard if it is wrong.

Real-Resource VerifyPrice SLOs (Constitutional):

22.7 VerifyPrice in Practice

At the North-Star level, VerifyPrice answers a simple allocator question:

> *"If I hold this asset for a cycle, does one unit still buy at least as much verification as it used to?"*

VerifyPrice matters in three dimensions, designed to avoid circular reasoning:

22.7.1 The Constitutional SLO Band

These are **non-negotiable** targets, measured in real resources on reference hardware. They are **exogenous to token price**. Whether the token rallies or dumps, these targets must hold. If verification takes 30s on a laptop, the "anyone can verify" promise is broken. The table below is the one referenced as Section 22.7's SLO band; the measurement specification that produces its inputs is Section 22.6.

What "Exogenous" Does and Does Not Cover

Exogeneity to token price is the property needed here, and it is the one being claimed: the target must not be redefined because the market moved.

It is not exogeneity to everything. Measuring in real resources on reference hardware makes these SLOs dependent on two conditions a state can adjust—whether unprivileged reference hardware remains obtainable, and what power costs at the edge and at the prover. Both are policy surfaces (Section 5.5).

The consequence is that Red Line 1, the thesis's hinge, is externally triggerable. This is monitored upstream via sovereign optionality (Section 17.8) and Red Line 13 rather than left to be discovered at the hinge. Section 17.10.1 separates the verification-side exposure (hardware obtainability) from the proving-side exposure (bulk energy), which behave differently and should not be conflated.

There is also a reason the exogeneity requirement runs deeper than engineering preference. If the framing of Section 2.0.1 is right that a medium restructures only those who *habituate* to it, then verification must be cheap enough to become a habit before it can

Workload	Modality	Target t_{95} (Laptop)	Sev-1	Prove-side feasibility
`PROOF_2^20`	M1	$\leq$ 5s	$>$ 10s for 7 days	Production; seconds to minutes on Server-Class
`MATMUL_4096_FP32_FREIVALDS`	M2	$\leq$ 5s (I/O-bound: 192 MB read and hashed)	$>$ 10s for 7 days	Production; one pinned-kernel product
`MATMUL_4096_FP32_SNARK`	M1	$\leq$ 2s	*registry target; not constitutional until the maturity gate lifts*	Research; no production prover at buyer-acceptable cost
`INFER_LM_7B_512TOK`	M3 (sampled re-execution)	$\leq$ 30s on **Server-Class** (registry-bound); Laptop-Class re-execution of a 7B forward pass over 512 tokens is minutes and is published, not targeted	$>$ 60s Server-Class for 7 days	Production; the prover runs the model once
`INFER_LM_7B_512TOK_ZKML`	M1	$\leq$ 2s (once wrapped, proof check is milliseconds, independent of model size)	*registry target; not constitutional until the maturity gate lifts*	Research; peer-reviewed ZKML proves a 7B forward pass over 2,048 tokens in about ten minutes on one GPU [Sun et al. 2024]
`INFER_LM_70B_256TOK`	M3	Server-Class only (model exceeds Laptop-Class memory); *aspirational as an M1 SKU*	n/a until ZKML matures	M3: production on Server-Class; M1: not currently achievable
All	—	failure rate $\leq 0.1\%$	$>$ 1% for 7 days	—

Table 22.10: Constitutional VerifyPrice targets, banded by modality. The band is not decoration: M1 verification, once the proof is wrapped, is milliseconds independent of model size, and M3 verification of a language model *is* running the language model, so a single "INFER_LM_7B $\leq$ 30s" row was either an M1 promise no prover can yet keep or an M3 promise no laptop can keep. Rows marked "registry target" are published so the destination is public; placing them in the constitutional set as though they were deliverable would trip Red Line 1 on publication day, and they become constitutional only when the maturity gate of Section 22.10 lifts for their component. Where the reference class is Server-Class, the Laptop-Class figure is still published.

become anything else. If its cost tracks the token's price, habit never forms, and neither does whatever the habit would have built. The SLO is not a performance target; it is the precondition of the medium having any effect at all.

22.7.2 Protocol Affordability SLOs (Operational)

These measure whether verification remains affordable as a fraction of typical transaction costs:

$$\text{AffordabilityRatio}(W) = \frac{\text{VerifyCost}(W)}{\text{MedianFee}(W)}$$

- **Target:** AffordabilityRatio $\leq 5\%$ for all workload classes.
- **Settlement:** Verification cost $\leq 1\%$ of median transaction value.

Read the denominator before celebrating the ratio. The ratio improves in two ways: verification gets cheaper, or the fee gets bigger. A protocol that raises MedianFee(W) passes this SLO while becoming less affordable to use, so the ratio must never be published without its components, and a falling ratio driven by a rising denominator is a deterioration wearing an improvement's clothes. The settlement variant has the mirrored exposure: if median transaction value is measured in token terms, appreciation mechanically lowers the ratio and re-imports token price through the back door — the exact circularity the physical SLOs exist to exclude. Three guards hold the family honest, each of a different kind. The first is component disclosure: the fee-based ratio is always published alongside the fee level itself, so a denominator-driven improvement is visible on inspection. The second is a unit convention: the settlement ratio is computed on USD-referenced transaction value, never token-referenced. The third is a monotonicity rule: the affordability SLO is credited only when the ratio falls while the fee level is flat or falling, and improvements achieved on a rising fee are reported as neutral, not as passes. The affordability family is operational, not constitutional, and these guards are why it stays that way honestly.

22.7.3 Token-Quoted VerifyPrice (Market Signal, Not Target)

Token-quoted VerifyPrice is a **useful market signal**, but not a constitutional target. For a scalar reading, fix a single reference workload W^* and take the cost-percentile field only, with the fee schedule expressed as a fixed token quantity:

$$\text{VerifyPower}(\text{token}; W^*) = \frac{q_{\text{fee}}(W^*) \cdot P_X}{p_{50,c}(W^*)}$$

where $q_{\text{fee}}(W^*)$ is the fee schedule for W^* denominated in units of the native asset X, and P_X is the asset's USD reference price at the observation — so the signal carries the token price explicitly on the right-hand side, which is what makes it a *market* signal rather than a disguised operational one. A vector has no reciprocal; the scalar must be constructed by naming the workload and the field. Which workload is the reference is a publication choice, not a free parameter — it is declared, pinned, and changed only with notice, so the signal cannot be dressed by re-picking W^*. Note the family relationship: this quantity is the token-side image of the AffordabilityRatio, and the two must be published together so neither is read alone; the affordability ratio is the operational SLO, this is its purchasing-power translation.

Token price is endogenous. Making this a *target* creates circular reasoning. The protocol commits to **real-resource SLOs and affordability ratios**. Token-quoted metrics are dashboards for market participants, not governance constraints.

Level	What It Measures	Who Enforces	Con
Real-resource (Constitutional)	Verification cost in real resources on reference hardware	Protocol governance	Sev-
Affordability (Operational)	Verification as % of fees	Fee policy	Fee
Token-quoted (Market)	Purchasing power signal	Market participants	Info

SLO Hierarchy:

22.8 Physical VerifyPrice: Auditing the Infrastructure Behind a Workload

Note on terminology: Section 22.7 names the constitutional cost band on verifying a *proof or compute claim* in real resources on reference hardware the *Real-Resource VerifyPrice SLO*. This section introduces the metric that carries the name "Physical VerifyPrice" alone: the cost of verifying the *physical infrastructure claims* (energy, hardware, jurisdiction) behind that proof or compute claim, i.e., the Facility Capacity Receipts introduced in Section 17.7.

VerifyPrice tells us whether a proof or compute claim is cheap to check. **Physical VerifyPrice** tells us whether the physical claims behind that proof are cheap to audit. A Work Credit cannot be high-assurance if verifying the proof takes seconds but verifying the facility, hardware profile, energy receipt, and jurisdictional risk takes months of trust-the-vendor diligence.

Concretely, Physical VerifyPrice asks:

- Can an allocator verify the facility's claimed energy profile?
- Can an auditor verify outage and curtailment history?
- Can the network verify that the hardware profile was not silently substituted?
- Can a user verify that compute was not routed through a sanctioned or coerced jurisdiction?

Physical VerifyPrice

The time, cost, and confidence required to verify the physical infrastructure claims (energy, hardware, jurisdiction) behind a unit of verified work.

For each facility profile, the stack should publish p50/p95 time and cost to verify FCR claims, sampling frequency, third-party challenge success rates, and unresolved discrepancy counts. If Physical VerifyPrice rises above threshold, credits from that facility receive reduced issuance weights or issuance caps (Section 25.6.5). This metric is formalized alongside the VerifyPrice model in Appendix A.

22.9 Layer-4 Stress Tests

Layer 4 passes its SoV audition if it survives several adversarial scenarios:
Circuit bloat. Can we detect when proofs become too expensive to verify? Is there a migration path to leaner circuits?
Prover cartel. Can neutral routers and SLA/slashing mechanisms prevent a small set of proof factories from monopolizing high-value workloads?
Proof system break / new attack. Can we deprecate a proof system, rotate to alternatives, and quarantine affected Work Credits?
Delegation depth (joint completion). The VERIFYPRICE SLO reads on one verification. A delegated task is a tree of them (Section 7.2), and a tree's tail is not the tail of any one of its hops. Does a deep tree still complete inside its refund window?

Why the single-verify SLO does not cover a tree. Take a task that fans out to k children at each of d levels, every hop returning a receipt its parent verifies. The tree has $N = (k^{d+1} - k)/(k - 1)$ hops; a parent cannot accept until its slowest child has returned, and the levels run in sequence. Two consequences follow from arithmetic alone. First, the tree's completion time is governed by the *far* tail of the single-verification distribution, not by its p95: with hops treated as independent, the chance that all N land inside their individual p_{95} is 0.95^N, which for the reference tree below is of order 10^{-8}. What matters is the p99.9 and beyond, which the canonical tuple does not publish. Second, failure compounds: at the reference failure SLO of 0.1% per verification, a 340-hop tree completes with no failed hop about 71% of the time, so roughly three trees in ten exercise a retry or a refund path somewhere. Independence is the favorable assumption. Under the common-cause scenarios of Section 17.9 hops fail together, and the tree aborts more often than the product says.

Scenario. For a published reference tree—$(d, k) = (4, 4)$, 256 leaves and 340 hops, chosen as a first reference and not as a threshold—report, for each canonical workload and under each

Layer 0 disruption scenario already specified: the p50 and p95 completion time of the whole tree on Laptop-class verifiers; the probability that the tree completes before its outermost refund window closes; and the hop or settlement edge whose failure aborts the most leaves. Publish the three beside t_{95} on the Proof & Compute Board. No SLO is set on any of them in this version; the purpose is to find out whether the single-verify SLO is the binding one for the workloads that delegation would generate, before anyone assumes it is.

If Layer 4 remains **cheap to verify, open to participate, and instrumented enough to handle change**, then Proofs and Compute deservedly move closer to "monetary primitive" rather than "platform feature."

22.10 Minimum Viable Layer-4 Economy

Canonical Workload Starter Set: The starter set is the subset of the registry (Table 22.4) that a minimum viable Layer-4 economy launches with. Identifiers are the registry's; the columns here are the ones a market maker reads.

WorkloadID	Description	Modality	Commercial target	WC tier ceiling
`PROOF_2^20`	Generic ZK proof	M1	$t_{95} \leq 5$s	A (full)
`MATMUL_4096_FP32_FREIVALDS`	Matrix multiply 4096×4096, Freivalds-audited	M2	$t_{95} \leq 5$s	B (discounted)
`MATMUL_4096_FP32_SNARK`	Matrix multiply 4096×4096, succinctly proved	M1	$t_{95} \leq 2$s	A (gated)
`INFER_LM_7B_512TOK`	7B-parameter inference, sampled re-execution	M3	$t_{95} \leq 30$s (Server-Class)	B (discounted)
`PROVENANCE_MEDIA_V1`	Media provenance chain	M1	$t_{95} \leq 3$s	A (full)
`SETTLE_ATOMIC_BTC_XMR_V1`	Atomic swap settlement, BTC↔XMR	M1 receipt over chain finality	$t_{95} \leq 5$s	A (full)

Table 22.11: Canonical workload starter set, drawn from Table 22.4. Two target levels are distinct by design and must not be read as one: the constitutional floor (Section 22.6 SLO table) is the minimum a network must hold to remain constitutional, while this starter-set target is the commercial bar the registry holds itself to for the named SKU. Where the two coincide, as they now do for the inference row, it is because the constitutional table was rebanded by modality and the earlier "$\leq$ 10s" commercial figure for a 7B re-execution on Laptop-Class was not a target any verifier could meet. WC tiers shown are the *ceiling* each workload can reach once its verification component is Production-mature; current effective tier is the lower of the ceiling and the maturity-gated tier (Section 22.10 tier rules). The `_SNARK` SKU is listed so the M1 destination for matrix work is public; it issues nothing until a production prover exists.

Maturity gate. No workload may issue Tier A (full-weight, pristine-collateral) Work Credits while the maturity label of its verification component (Table 22.5) is Experimental or Pilot. Until a component reaches Production, its workloads issue at Tier B or below regardless of

the mathematical tier their proof type would support. As of this writing that gate is binding on every PoUW-based row in the starter set: MatMul-PoUW *consensus security* is a conjecture (the verification asymmetry itself is Freivalds' theorem [Freivalds 1977] and is not what the gate reads on; Section 22.4.1), succinct MatMul proving is research, inference PoUW is pilot-stage, and pricing pristine collateral over any of them would define a AAA rating on an instrument whose proof system has not survived production. For starter-set rows whose verification component does not appear in Table 22.5 at all (`PROOF_2^20`, `PROVENANCE_MEDIA_V1`, `SETTLE_ATOMIC_BTC_XMR_V1`), the gate reads on the maturity of the underlying proof system family — generic succinct proofs, provenance attestation, atomic-swap verification — each of which the maturity table tracks at the family level; a family not yet tracked is treated as Pilot until it is. The gate lifts only when the maturity table lifts, and the maturity table is moved by deployment evidence, not by argument.

Tier	Verification Type	WC Weight	Collateral Grade
Tier A	Full cryptographic proof (ZK-SNARK/STARK) *and* component at Production maturity	1.0×	Pristine
Tier B	Probabilistic verification (audited transcripts)	0.6×	Standard
Tier C	Attestation-backed (TEE + sampling)	0.3× (service-grade claims only)	Ineligible

Table 22.12: Work Credit issuance tiers.

Tier Rules for Work Credit Issuance:

Fee + Burn + Slashing Logic:

- 70% of fee to capacity providers and their delegating stakers (reward)
- 20% retired (a pro-rata buyback; net-issuance arithmetic, not moneyness)
- 10% to the assurance budget (Layer 0 sampling, audits, security review)

This is the reference-design partition pinned in Section 25.6. Its 20% burn share is a deliberately conservative variant of the 30–50% band in Chapter 8: the reference design also funds an assurance budget, so the three shares are set to sum to 100%. The band remains the design space; a 20% burn is the harder case for the value-capture claim, not the easier one.

Fee formula (reference):

$$F(W) = \text{BaseFee}(W) \times (1 + \text{CongestionMultiplier}) \times \text{SLATierMultiplier}$$

Chapter 23

Layer 5: Value & Settlement

Layer 4 turns work into truth. Layer 5 turns truth into **money movement**.

It is the layer where balances change, salaries get paid, cross-asset swaps execute, and collateral gets posted or released.

The constraints are harsh:

Repression-resilience. Default assumption: yield-curve control, capital controls, KYC choke-points.

Privacy & lawfulness. Users need *privacy by default* and *auditability by consent*.

Non-custodial safety. Settlement protocols must be self-custodial and refund-safe.

Privacy Rails and Lawful Obligations

Privacy rails are designed for lawful confidentiality and user-controlled disclosure; they are not framed as tools for evading legal obligations. Some regulatory demands may be incompatible with privacy-preserving systems; in those cases, the protocol acknowledges the tension rather than promising universal compliance. See Chapter 27 for the full legal posture.

23.1 Settlement as a First-Class Workload

The stack treats settlement itself as a **canonical workload**:

- Workload type `Settlement(S, policy)` includes chain(s) involved, transfer size, corridor routing, and policy hooks.
- Proofs at Layer 4 can attest that settlement followed policy and the final state matches expectations.

This means:

- Layer 5 is not "just payments"; it is a **proven computation** over state transitions.
- Settlement events generate PIDL receipts that can be audited, archived, or collateralized.

23.2 Design Constraints for Privacy Rails

Privacy rails must satisfy a tight triangle:

1. **Non-custodial.** Users never relinquish both principal and unilateral control to intermediaries. Adaptor signatures, HTLC-like constructs, or scriptless scripts enforce this.
2. **Refund-safe.** If a swap fails, no party loses value, and every party has a unilateral no-loss exit within bounded time — an exit that may be a refund of its own leg or a claim on the counterparty's, depending on which side of the protocol it stands on (Section 23.4). Refund logic is testable and provable.
3. **Privacy-by-default with optional disclosure.** Default flows have strong anonymity sets; viewing keys allow specific flows to be revealed.

This is where the **Privacy Rails Kit (PRK)** lives.

Policy = Predicates, Not Surveillance

Policy compliance in Layer 5 is achieved through **predicate proofs**, not graph inspection.

What can be proved (ZK-compatible):

- Not on sanctions list → Membership proof: "sender $\in$ AML-cleared set S"
- Within limits → Range proof: "amount $\leq$ \$10,000"
- Authorized counterparty → Set membership proof
- Tax reporting → Viewing key + receipt disclosure

What can be proved about one's own inputs only:

- "None of *my* inputs to this transaction is in published exclusion set E" — provable as non-membership against a committed set, in the manner of Privacy-Pools association sets [Buterin et al. 2023]. The proof reaches exactly one hop: it says nothing about where those inputs came from before the prover held them.

What cannot be proved, and is not attempted:

- "Never interacted, at any distance, with blacklisted address X" — taint at distance requires global graph knowledge, and the protocol does not have it
- "Full transaction history disclosure" — defeats privacy model
- "Real-time surveillance feed" — centralizes and doxxes

23.3 The Privacy Rails Kit (PRK)

PRK is the Layer-5 sibling of PAL. It gives builders a way to express **pay-for-proof** and **pay-for-compute** intents and execute them over privacy-preserving, non-custodial corridors.

At a high level, PRK supports:

- **BTC↔XMR corridors, and a gated BTC↔ZEC corridor.** Using adaptor signatures, ensuring atomicity. Prior art is directly deployed *for BTC↔XMR only*: the Bitcoin–Monero atomic-

swap protocol of Gugger (2020) and its production implementations (COMIT's `xmr-btc-swap` [COMIT 2021], UnstoppableSwap [UnstoppableSwap 2025]) are the operative reference, and Section 23.4's refund-safety invariants are stated so that such corridors can be certified against them rather than merely named. BTC↔shielded-ZEC adaptor-signature swaps **do not exist in production**. Transparent-ZEC HTLCs exist and forfeit privacy on the ZEC leg, which is the leg the corridor exists to protect; a shielded corridor requires scripting or adaptor constructions inside the shielded pool that Zcash does not currently offer [Hopwood et al. 2025]. The ZEC corridor is therefore registered (Table 22.4) and gated on that work: it is a design target, not a route, and nothing in this document counts it toward a minimum viable stack.

- **Shielded pools and internal flows.** For in-asset privacy with proofs for policy compliance. Outflows from each pool cannot exceed the value that entered it.
- **Conditional payment flows.** "Pay address A if proof P of workload W is posted within time T; else refund."

  PRK exposes intents like:

```
pay_for_proof(claim_id, max_price, corridor_policy)
execute_corridor(btc_input, xmr_output, privacy_policy)
batch_payroll(payroll_blob, corridor_set, audit_policy)
```

23.4 Settlement Safety: Refund and Bridge Invariants

Non-custodial rhetoric is cheap; invariants are not.

Refund Safety Invariant: Refund safety is often stated as "both sides can unilaterally reclaim funds," and the deployed prior art does not give that. The Gugger (2020) BTC↔XMR protocol gives *no-loss-or-punish*, which is asymmetric by construction: Monero has no scripting, so the XMR leg cannot time out on its own. If the BTC-side party goes offline past the refund window, the XMR is unrecoverable by either side and the protocol's remedy is the *punish* path, which delivers the BTC to the XMR-side party. Nobody ends the swap poorer than they began it; the exits are not symmetric refunds. The invariant is therefore stated in the terms the protocol actually satisfies:

Refund Safety

A corridor C is **refund-safe** if, for every transaction T and every party P to it, P has a *unilateral no-loss exit* available within a bounded time τ_C from the moment T leaves the happy path — an action P can take without the counterparty's cooperation after which P holds value at least equal to what it committed (its own leg refunded, or the

counterparty's leg claimed). Formally,

$$\text{refundSafe}(T) = 1 \iff \big(T \text{ succeeds}\big) \ \vee \ \big(\forall P \in T:\ P \text{ exits without loss within } \tau_C\big).$$

Where:

- τ_C is bounded and published per corridor (e.g., ≤ 24 hours for BTC↔XMR at current timelocks) and is the sum of the longest refund window and the longest confirmation delay on either chain. The XMR leg's confirmation depth is set above 18 blocks, the depth of the September 2025 reorganization [Rucknium 2025].
- An exit is unilateral if it requires no counterparty signature after the setup phase. Which exit a party holds (refund or punish) depends on which leg it stands on and is published in the corridor's profile.
- Success means both legs complete atomically.
- The metric refund_safe(C) of Section 23.5 is computed over *protocol-attributable* failures: a party that has an available no-loss exit and fails to take it within τ_C (client offline, keys lost) is recorded under a root-cause class and does not count against the corridor, because a hard 1.0 that trips on user liveness is a metric no corridor with humans on it can hold.

Bridge Safety: Cross-chain bridges were, in 2021–22, the single biggest loss vector in the industry. Layer 5 imposes:

- **No unclassified signer sets.** A small permissioned signer set is admissible only as Class B or Class C at the haircuts below, disclosed as such in the profile, and never presented as trustless. A slogan such as "no magic multisigs" beside a table that admits a 5-of-9 at a 10% haircut is rhetoric the table contradicts; the classification replaces the slogan, and the haircuts below are set to match it.
- **Light-client or proof-based validation where possible.**
- **Explicit trust classifications** reflected in profiles and telemetry.

Bridge Class	Trust Model	Collateral Eligibility	Example
Class A: Trustless	Cryptographic only: light-client or ZK-verified state proof of the source chain	Full	Light-client / ZK-verified bridge
Class B: Threshold	m-of-n MPC or multisig, signers permissioned	30% haircut	5-of-9 threshold signer set
Class C: Federated	Known federation of functionaries; no on-chain bond	50% haircut	Liquid-style
Class D: Custodial	Single custodian	**Not eligible**	CEX-wrapped

Table 23.1: Bridge trust classification. Liquid's federation is a set of named functionaries running hardware signing modules; it is not bonded in the sense of posting slashable collateral, and the earlier "bonded" label is withdrawn.

An atomic swap is not a bridge — no asset is represented on a foreign chain, so there is nothing to classify — and the PRK corridors of Section 23.3 are classified as corridors, not under this table. They carry Class-A-equivalent trust (cryptographic only), which is the reason the kit prefers a corridor to a bridge wherever both exist for a pair.

Rule: Class D custodial bridges are excluded from SoV backing. Assets held via Class D bridges do not count toward collateral or reserve calculations.

23.5 VerifySettle: Settlement Telemetry

VerifySettle tracks:

- **Success rates** across corridors (p50/p95, by size bucket).
- **Refund outcomes** — how often refunds are exercised and their timing.
- **Anonymity-set health** — shielded pool sizes, churn, volume.
- **Concentration metrics** — top-N LPs, corridor operators, jurisdictional spread.

Corridor Health SLOs:

Anonymity is measured per asset, because the assets hide differently. A single "anon_set ≥ 1000" was unreachable for one leg and trivial for the other. Monero's anonymity is per-output: a ring of 16 today (FCMP++, which would make the effective set the whole output set, is on testnet with no mainnet activation date set as of this writing), so the meaningful metric is ring size times decoy quality — the fraction of rings whose decoys are indistinguishable from the real spend under the published heuristic set. Zcash's anonymity is pool-wide: the shielded pool's size times its spend churn over the window, so a pool of a million notes with no spends is a museum, not an anonymity set. The thresholds are:

- **XMR leg:** ring size at protocol minimum (16) with $\geq 95\%$ of corridor rings passing the decoy-quality heuristic; Sev-1 if decoy quality falls below 80% for 30 days. A "≥ 1000" set is not a Monero quantity until FCMP++ ships, at which point the metric is revised to effective set size.
- **ZEC leg (when the corridor is admissible):** shielded-pool spend churn ≥ 1000 distinct shielded spends per day on the pool the corridor uses, with corridor volume $\leq 10\%$ of that churn so corridor notes do not stand out; Sev-1 if churn falls below 500 for 30 days.
- **BTC leg:** the btc_link metric above; Bitcoin has no anonymity set and the corridor's job is not to make its BTC outputs distinguishable from ordinary spends.

23.6 Layer-5 Stress Tests

Layer 5 earns the label "Value & Settlement" if it can navigate:

Metric	Definition	Target	Sev-1
swap_success(C)	% of swaps that complete	$\geq 95\%$	$< 90\%$ for 7 days
refund_safe(C)	% of off-happy-path transactions (those that did not complete atomically) in which every party had a no-loss exit within τ_C (Section 23.4). The complement is the *protocol-attributable* failure rate. Failures caused by a party's own environment (client offline past the window, keys lost) are published under a root-cause class and excluded from the denominator; the *overall* figure that counts them is published beside it and is what the Phase I gate and Red Line 2 read at 99.5%	$\geq 99.9\%$ (protocol denominator)	Any protocol-attributable loss
ttf(C)	Time-to-finality (p95), bounded below by confirmation depth on both chains — 30–90 minutes on the deployed BTC↔XMR corridor	≤ 90 min	> 3 hours
anon_set(C)	Per-asset anonymity metric, defined per leg (below)	Per-asset thresholds below	Per-asset thresholds below
btc_link(C)	BTC-leg linkability: share of corridor BTC outputs that a published heuristic set (common-input, change, timing, amount) links to a prior corridor output	$\leq 20\%$	$> 50\%$ for 30 days
lp_concentration(C)	Top-3 LP share	$< 50\%$	$> 70\%$ for 14 days

Table 23.2: VerifySettle corridor health SLOs.

On-/off-ramp strangling. Can users still move value between BTC and privacy assets without centralized exchanges in country C? Do non-custodial corridors maintain $\geq 95\%$ success with no protocol-attributable loss?
Corridor LP exit. If a major liquidity provider disappears, do corridors collapse, or does routing adapt?
Regulatory attack on privacy. Can lawful-privacy corridors (viewing keys + receipts) keep enough demand to sustain anonymity sets when shielded assets are labeled "high risk"?

23.7 Minimum Viable Layer-5 Corridor

Disclosure Scope Model: This is easily mistaken for a viewing-key scope model, and a viewing key cannot deliver the four scopes it requires. A shared viewing key is knowledge: once given it cannot be revoked, it reads every note the address ever receives rather than a time window, and it reveals every amount rather than only those above a threshold; and hierarchical derivation, whatever it is called, has a root. The scopes are real requirements, and each is met by a different mechanism — only one of which is a viewing key.

Scope	What Is Revealed	Mechanism	Use Case
Payment-scoped	One payment's amount, memo, and recipient	Per-note payment disclosure: the sender reveals that note's decryption material (a payment-disclosure proof), not an address key	Receipt for a specific payment
Epoch-scoped	All receipts to one address family within an epoch	Epoch-rotated addresses: the user derives a fresh receiving address per epoch and hands the auditor the incoming viewing key for the epochs in the mandate only. Scope is achieved by rotation, not by a time-bounded key	Tax audit over a fiscal period
Threshold-scoped	That no payment in a set exceeded (or that some payment exceeded) a threshold, without amounts	ZK selective-disclosure predicate: a range or aggregate proof over the note set, checked by the compliance party; no key changes hands	AML / limits monitoring
Full	All transactions to a spending key's addresses	The user's own full viewing key, never shared	Personal records, self-audit

Table 23.3: Disclosure scope model, with the mechanism that delivers each scope. A scope not on this list is not offered, and a viewing key is the mechanism for exactly one row.

Key management: Users generate and control all disclosure material. There is no revocation of a shared viewing key — it is knowledge, and knowledge is not recalled — so the operative discipline is *forward rotation*: receiving addresses and their incoming viewing keys

rotate per epoch, a disclosed key exposes past epochs only, and future epochs are dark to every past auditee by construction. Derivation is hierarchical from a spending-key root the user holds; that root *is* a master key for its own tree, which is why it is never disclosed and why per-epoch keys, not the root, are the unit of disclosure.

Requirement	**Admission threshold (entry to default routing)**
refund_safe	No protocol-attributable loss in the certification suite and in $\geq$ 90 days of monitored operation; $\geq$ 99.9% thereafter (the Sev-1 condition is any protocol-attributable loss, a stricter reading of the same series)
swap_success	$\geq$ 95% over the trailing 30 days (Sev-1 is $<$ 90% for 7 days)
ttf (p95)	$\leq$ 90 minutes (Sev-1 is $>$ 3 hours)
anon_set	Per-asset admission bar of Section 23.5: XMR ring at protocol minimum with decoy quality $\geq$ 95%; ZEC churn $\geq$ 1000/day when that corridor is admissible; btc_link $\leq$ 20%
Settlement type	Corridor (atomic swap), or bridge of Class A–C at the published haircut; no Class D leg
Disclosure support	Payment-scoped disclosure minimum (Section 23.7)

Table 23.4: Minimum viable corridor requirements. Admission thresholds are deliberately tighter than the Sev-1 thresholds of Section 23.5: a corridor enters default routing above the target band and is removed from it at the Sev-1 line, so that the two events are separated by a margin rather than a coin flip. Setting admission equal to Sev-1 would list corridors at the edge of removal.

Minimum Viable Corridor Spec: Corridors failing the admission thresholds are flagged as "experimental" and excluded from default PRK routing; a listed corridor that crosses a Sev-1 line is handled by the incident path of Section 25.5.3.

Chapter 24

The Modular Stack

Builders touch **modules**, and underneath each module the system enforces **primitives**. Parts III and IV described the stack layer by layer; this chapter turns the view to modules and primitives.

24.1 Primitive Catalog: The Twelve Pieces

Twelve primitives, and collectively they do one thing: they let two strangers trade verified work without either one trusting the other's platform. That is the whole catalogue. When proofs from this kit clear, the network mints verified FLOPs rather than heat; when the receipt schema binds, a claim becomes portable instead of platform-bound; when the corridors settle non-custodially, the value moves without a custodian appearing in the middle. No single primitive accomplishes any of that. The twelve are what the stack ships so that trust can be withdrawn, one counterparty at a time.

They fall into four clusters:

Compute / Consensus Primitives (AI-Money Substrate)

1. **MatMul-PoUW kit** — A work function that turns matrix multiplication into a useful lottery with verification asymmetry. Canonical sizes (the `MATMUL_4096_FP32` SKUs of Table 22.4), seeding rules, and the adversarial acceptance suite ensure that:
   - production cost scales $\sim O(n^3)$,
   - verification cost scales $\sim O(n^2)$ [Freivalds 1977] (Freivalds; a theorem),
   - and cheating is more expensive than being honest — which, for the consensus role, is the conjecture of Table 22.5 and the open problem of Section 22.4.1.

   This is the Duplex-style substrate: when proofs from this kit clear, the network mints **verified FLOPs** rather than heat.
2. **Verified-inference harness** — An Ambient-style proof-of-logits layer [Ambient 2024] for AI workloads: deterministic transcripts, randomized audits, peer-prediction, and stake-and-slash for dishonesty. Where full ZKML is not yet practical, the harness provides:
   - probabilistic guarantees that outputs came from model M under policy P;
   - receipts per inference that VCO and PFS can route and price.

As ZKML matures, the same harness becomes the "front door" for full proofs. Until then, it is what keeps an inference receipt from being a screenshot of a trust-me box: a transcript exists, an auditor can re-run it, and a prover who fabricates one loses stake.

3. **Canonical workload registry** — The dictionary of standardized SKUs for useful work, published as Table 22.4: `PROOF_2^20`, the two `MATMUL_4096_FP32` SKUs, `INFER_LM_7B_512TOK`, `INFER_LM_70B_256TOK`, `PROVENANCE_MEDIA_V1`, `SETTLE_ATOMIC_BTC_XMR_V1`, and the gated ZEC corridor. This is the thesis's unit of account in physical form: without it, "compute" is a negotiating position rather than a quantity, and no two quotes mean the same thing.

Inference Proof Tier Taxonomy Verified inference is the most economically significant—and currently most fragile—market under the "AI Money" analytical lens. The label does not make inference claims monetary; the guarantees are tiered explicitly so weak assurance is not marketed as strong:

Tier	Verification Type	Description	Error Bound	WC Weight
Tier A	Cryptographic correctness	Full ZK proof of inference (ZKML)	0 relative to the proved (quantized) computation; quantization gap to the advertised FP model published separately	1.0×
Tier B	Probabilistic soundness	Audited transcripts + randomized checks + peer prediction	Explicit (e.g., 99% confidence)	0.6×
Tier C	Attestation-backed	TEE attestation + sampling audits	Implicit (trust TEE + sampling; attestation not relied on against a prover with physical access)	0.3×

Table 24.1: Inference proof tier taxonomy.

Use Case	Minimum Tier	Rationale
Work Credit issuance (full weight)	Tier A	High-assurance service claims require cryptographic verification
Work Credit issuance (discounted)	Tier B	Acceptable with explicit error bounds and service haircut
Work Credit issuance (service-grade, 0.3×)	Tier C	Can sell inference and mint service-grade Work Credits at 0.3×; cannot mint high-assurance Work Credits
Collateral (pristine)	Tier A only	DeFi integrations require cryptographic certainty
Collateral (standard)	Tier A or B	With appropriate haircut
Collateral (ineligible)	Tier C	Not acceptable as collateral

Table 24.2: Tier eligibility rules.

Why this matters: Without this taxonomy, "verified inference" conflates cryptographic proof with probabilistic audit with vendor attestation. By making tiers explicit and tying them to Work Credit issuance weights, the thesis acknowledges engineering reality while preserving

the instrument hierarchy: assurance grades service claims; monetary candidacy belongs only to the separate base asset.

Tier B requirements (reference): For Tier B verification to qualify:

- **Transcript reproducibility against a pinned kernel:** FP32 arithmetic is non-associative and BLAS kernels differ across hardware profiles, so bit-identical transcripts across heterogeneous provers are not available by default. The requirement is *pinned-kernel* reproducibility: each workload registry entry names a reference kernel (arithmetic order, FMA contraction off, no TF32), a transcript is reproducible by any prover running that kernel, and cross-kernel divergence is handled through the workload's relative error bound rather than exact-match. A prover may run a faster kernel only by proving conformance within ε_r of the reference result.
- **Audit rate:** $\geq 10\%$ of inferences are independently re-executed by random auditors.
- **Peer prediction:** Multiple independent provers; divergence beyond ε_r triggers investigation.
- **Stake-and-slash:** Provers stake collateral; fraudulent transcripts $\rightarrow$ 100% slash.
- **Error bound publication:** Explicit statistical guarantee (e.g., "99% confidence that output matches claimed model within ε_r").

Migration path: As ZKML matures: (1) Tier A coverage expands to more model classes; (2) Tier B issuance weight increases toward $1.0\times$ as probabilistic guarantees tighten; (3) Tier C's $0.3\times$ weight is stepped toward zero and Tier C is eventually retired from WC eligibility. The goal is for inference verification to converge to Tier A over time, but the system functions with tiered guarantees in the interim.

Proof Primitives (Receipts as First-Class Objects)

4. **PAL SDK (Proofs-as-a-Library)** — Developer surface for Layer 4.
5. **Multi-ZK adapter & auto-selector** — Policy engine keeping proofs a commodity, not a vendor feature.
6. **PIDL (Proof Interface Definition Language)** — Minimal receipt schema binding claim, workload, proof, SLA, timestamps, and signatures.

Privacy & Settlement Primitives (Private-Money Substrate)

7. **Adaptor-signature atomic-swap kit** — Library for non-custodial BTC↔XMR settlement with refund-safety as defined in Section 23.4, and the shielded-scripting work that a BTC↔ZEC corridor would need before it can be built. The adversary here is the counterparty who takes your funds and disappears mid-trade, and the kit exists so that no one—including the corridor operator—can be that counterparty. Non-custodial rhetoric is cheap; invariants are not.
8. **Lawful-privacy corridor pack** — Hooks for viewing keys + receipts enabling regulated

actors to prove compliance. A bank that cannot demonstrate compliance to its regulator cannot route anything, so these hooks are what keep lawful privacy lawful rather than merely private.

9. **Bridge-safety templates** — Pattern book for cross-chain settlement that classifies every signer set and haircuts it (Section 23.4). The graveyard this pattern book walks past is full of bridges whose security model was, on inspection, five signatures held by people who knew each other — and whose documentation did not say so.

Market & Telemetry Primitives

10. **SLA escrow & slashing contracts** — Standard contracts for Bronze/Silver/Gold tiers with automatic refunds or slashes.
11. **VerifyPrice telemetry & methodology** — Measurement side of verification asymmetry.
12. **Neutral router & fairness tests** — Matching logic with house-share caps and entry-latency metrics.

PFS (Proof Factory Stack) bundles primitives 4–6, 10–12 to turn claims into routed proof jobs.

VCO (Verified Compute Orchestrator) bundles 1–3, 5, 10–12 to turn useful workloads into Work Credits.

PRK (Privacy Rails Kit) bundles primitives 7–9, with 6 for the receipt schema and 11 for VERIFYSETTLE telemetry, to turn a payment intent into a non-custodial, refund-safe, selectively disclosable settlement. Every primitive now belongs to at least one bundle.

24.2 The Four Reference Applications at a Glance

The four applications that follow are not a product roadmap. They exist to show that the stack closes: each one exercises a different span of layers, reports to a different set of public boards, and would trip a different set of red lines if the thesis is wrong. An architecture that cannot name which application would reveal which failure is not falsifiable in practice, however many metrics it publishes.

Between them the four applications stress eleven of the eighteen red lines. The seven they do not reach—Red Line 4 (telemetry capture), Red Line 8 (governance capture), Red Line 14 (the capturable wedge closing), Red Line 15 (the native collateral–capacity spiral), and Red Lines 16–18 (the three link-9 monetary tests: convenience-yield null, information-sensitivity reversion, denomination null)—are protocol-wide or market-wide failures rather than application-level ones, which is precisely why they cannot be caught by shipping an application and must be watched institutionally instead. Red Lines 16–18 read on holder behaviour and market structure of the base asset and cannot, by construction, be reached by any application. Red Lines 14 and 15 are the clearest cases: an individual application can be tech-

Application	Layers exercised	Reports to	Stresses	What its failure would show
Private treasury & payroll (Section 24.3)	L1 reach, L2 client integrity, L3 scoped claims, L4 policy proofs, L5 corridors	Settlement & Privacy; Agency Preservation	RL2, RL7, RL12	That lawful privacy is not operable at institutional scale—selective disclosure has become administrative disclosure
Media provenance (Section 24.4)	L0 attested capture, L2 distribution, L3 claims, L4 proof generation	Proof & Compute; Neutrality & Admission	RL1, RL3, RL10	That provenance cannot outlive the platform hosting it—receipts have become labels again
Verified inference (Section 24.5)	L0 capacity, L4 PoUW and proving, L5 contingent payment	Proof & Compute; Layer 0 Capacity; Economic Coverage	RL1, RL11, RL13	That verified compute earns no durable premium over ordinary cloud—the AI Money claim reduces to hosting
Proof/compute procurement (Section 24.6)	L4 workload registry, L5 settlement, L6 SLOs and telemetry; L3 principal binding in the agent-to-agent case	Economic Coverage; Value Capture; Wrapper Dominance; Agency Preservation (agent case)	RL5, RL6, RL9	That demand was speculative rather than budgeted, and value leaks to operators and wrappers instead of the asset

Table 24.3: The four reference applications mapped to layers, public boards (Section 26.9), and the red lines (Section 30.6) each would trip first.

nically successful while the capturable wedge closes or native collateral reflexivity degrades system-wide capacity.

24.3 Reference Application: Private Treasury & Payroll

It is the 26th of the month, and the payroll run is late. The company's CFO is in Singapore, its controller is in Lisbon, and its 340 employees are spread across eleven jurisdictions. Under the incumbent way, the 26th is late because payroll is a daisy chain of NDAs: the bank wants the salary list, the auditors want the bank, the tax authority wants the auditors, and every one of those disclosures is a copy of the transaction graph that will never be recalled. The controller cannot prove the run is correct without handing over the graph, and cannot protect the graph without becoming the only person who can prove anything. Every quarter the company chooses which of those two failures to have.

This application removes the choice. Salary amounts stay confidential, the transaction graph stays private, and the auditors still get scoped evidence that every payment was authorized, complete, and properly reported—not the underlying ledger, but proofs about it, limited to what the mandate actually asks.

The flow is the stack, end to end:

1. **Create.** Finance defines a payroll batch with policy.
2. **Compute.** Application logic computes net pays and corridor selection.
3. **Prove.** PAL generates proofs that payments are within policy.
4. **Settle.** PRK executes payments over BTC↔XMR corridors (the ZEC corridor when admissible).
5. **Verify.** Auditors verify correctness via proofs and receipts.

Design commitments. Policy conformance (each net pay within its band, batch total equal to the authorized total, every recipient in the authorized set) is proved under M1 as a succinct range-and-membership proof over the batch commitment, so the auditor's check is milliseconds and independent of headcount; settlement runs over the `SETTLE_ATOMIC_BTC_XMR_V1` corridor (Class-A-equivalent, no bridge), with the payroll disclosed to the auditor by epoch-scoped rotation and to each employee by payment-scoped disclosure (Section 23.7); and each PIDL receipt binds the batch commitment, the policy hash, the corridor identifier, both chains' finality proofs, the refund-safety outcome code, and the signatures of the finance approver and the proving factory.

24.4 Reference Application: Media Provenance

A clip surfaces on a Tuesday: a minister, a private conversation, forty seconds that would end a career if they were real. By Thursday the newsroom knows it is synthetic and cannot say so

with anything it can hold up. The platform's provenance label says "authentic"; the label is a string in the platform's own database, and everyone in the room understands what that is worth. The newsroom's correction cites "internal review." The clip keeps circulating. Nobody lied; there was simply nothing to check.

This application is what there is to check. Capture hardware attests at the moment of recording, provenance proofs travel with the footage rather than living in a platform's database, and the question "did this come from camera C and reach us through the stated chain?" acquires an answer a reader can verify without trusting the newsroom, the platform, or anyone's internal review.

The flow:

1. **Create.** Cameras capture footage with Layer-0 attested hardware.
2. **Prove.** Proof factories generate `Provenance(C, chain)` proofs.
3. **Settle.** Advertisers pay creators over privacy rails, conditioned on valid proofs.
4. **Verify.** End-users verify that content passed through the stated chain.

Design commitments. The provenance chain is the `PROVENANCE_MEDIA_V1` SKU, proved under M1 so a reader on Mobile-Class hardware verifies the whole chain in under a second; its trust root is a Layer-0 attested capture, which is M3 (hardware attestation) and is labelled as such in the receipt, so the modality of the chain's first hop is never overstated. Advertiser-to-creator settlement is a conditional payment over the BTC↔XMR corridor with no bridge in the path. The receipt carries the content hash, the capture-device attestation and its hardware profile, one signed edit record per hop with the hop's transform declared, the chain's proof, and the timestamps of capture and each hop.

24.5 Reference Application: Verified Inference

A risk officer at an insurer signs off on a model whose outputs will move payouts on thousands of policies. The vendor's API returns a number and a confidence score, and behind the endpoint the officer can see nothing: not whether the model is the one advertised, not whether the input was modified in transit, not whether the "confidence" is a measurement or a marketing decision. The incumbent answer is a due-diligence visit and an attestation letter. The officer files the letter and assumes residual risk, because there is nothing else to do with it.

This application replaces the letter with a receipt. The enterprise submits an inference request under stated constraints; the proof factory runs the model and produces `Inference(M, X; policy)` with a proof that the output came from model M on input X under that policy. Payment is contingent on the proof being valid, so the vendor is paid for verified work, not for uptime of a black box.

The flow:

1. **Create.** Client submits inference request with constraints.

2. **Compute.** Proof factory runs the inference.
3. **Prove.** PAL compiles into `Inference(M, X; policy)` with proof.
4. **Settle.** Payment via PRK contingent on valid proof.
5. **Verify.** Client verifies proof and accepts output.

Design commitments. The inference is the `INFER_LM_7B_512TOK` SKU under M3 (sampled re-execution, Tier B) today, with the `_ZKML` SKU as the M1 destination once a production prover exists; the officer's "proof" is, at present, a pinned-kernel transcript plus the audit record showing which random auditors re-executed it, and the receipt says so rather than calling it a proof. Payment is a conditional flow over the BTC↔XMR corridor, released against the receipt. The receipt binds the model hash, the input commitment, the output commitment, the pinned kernel ID, the transcript hash, the audit sample and each auditor's signed verdict, the assurance tier, and the prover's stake reference.

AI Money is an analytical lens or legacy label for these inference-service markets. The instruments are typed Work Credits and derivatives; they are not themselves money.

24.6 Reference Application: Proof/Compute Procurement

A DAO's treasury committee is staring at a proof-capacity crunch it can see coming and cannot buy ahead. Under the incumbent way, the committee's options are a cloud commitment letter—a price the vendor sets, hardware the committee cannot inspect, and a lock-in it will still be paying for after the crunch passes—or forward contracts on capacity nobody can verify exists. Neither instrument lets the DAO prepay for capacity it can later prove it received. The gap is not a pricing gap; it is a verification gap, and it is why procurement at distance has historically meant procurement on faith.

This application closes the gap. The DAO declares demand curves for canonical workloads over time, buys typed Work Credits as providers sell them against committed capacity, and holds them as prepaid service inventory. Anyone can audit the backing: the credits are reservations on registry workloads whose providers, proofs, and VerifyPrice profiles are public, so the committee's claim to future capacity is checkable by any token holder, not asserted by the committee.

The flow:

1. **Create.** DAO defines demand curves for workloads over time.
2. **Compute.** Procurement module computes optimal schedules.
3. **Prove.** As providers commit capacity, the DAO buys the credits.
4. **Settle.** Payments over privacy rails; typed Work Credits held as prepaid service inventory, not monetary reserves.
5. **Verify.** Anyone audits backing via VERIFYPRICE, profiles, and proof metadata.

Design commitments. The credits the DAO buys are typed by registry SKU and therefore

by modality: a `PROOF_2^20` credit is M1/Tier A and pledgeable, an `INFER_LM_7B_512TOK` credit is M3/Tier B and is prepaid inventory only, and the procurement module refuses to net the two into one "capacity" figure. Purchases settle over the BTC↔XMR corridor or, where the DAO's treasury sits on a Class A light-client bridge, over that bridge at full eligibility and never over Class D. Each Work Credit's PIDL receipt carries the SKU, the obligor, the reservation epoch, the issuance weight applied and why (tier, hardware grade, concentration), the facility's FCR reference, and the redemption terms, so the backing audit is a fold over receipts rather than a request to the committee.

The agent-to-agent case: delegation escrow. Replace the treasury committee with a software agent and the forward purchase with a task, and the same rail becomes the hop of Section 7.2. A specification enters: a registry workload, acceptance criteria, a deadline, an assurance tier. A receipt leaves. Between them sits the SLA escrow primitive of Section 24.1, with payment and the child's slashable bond both denominated in the base asset, released on a verifying receipt and refunded on timeout. Three commitments distinguish this from the committee's case. *Windows nest:* a parent's refund window must outlast the longest chain of its children's windows plus its own verification, or a parent can be slashed for a delay it could not have prevented; the joint-completion scenario of Section 22.9 is the test. *Escrow does not flow down the tree:* a parent is paid only on its own verifying receipt, so it funds its children's escrows from its own balance before it is paid. Working capital in the base asset therefore scales with the subtree an agent delegates, which is the operational reason a delegating agent holds a reserve at all (Section 7.5), and the reason a template shared by many such agents is a concentration worth measuring. *The principal is proven, not named:* the receipt may carry the optional principal_chain context (Appendix A) and the agent's principal binding (Section 20.5), which is what lets the Principal Map (Section 26.8.9) say whether the hop crossed a trust perimeter and whether a person can still revoke the agent that issued it. This is an extension of the present application and not a fifth one: it exercises Layer 3 where the committee's case does not, reports additionally to the Agency Preservation Board, and would show, if it failed, that delegation between strangers is not operable on receipts alone.

The DAO's treasury holds BTC/ETH reserve assets beside typed proof and compute capacity instruments. "ZK Money" and "AI Money" are analytical or legacy lenses for those service markets, not claims that the underlying Work Credits are money. Any separate base asset remains only a conditional monetary candidate.

24.7 Part IV Summary

Part IV is where the triad learns to **speak economics**:

- **Layer 4** turns computation and provenance into standardized, verifiable work units.
- **Layer 5** turns those work units into private, non-custodial value flows.

With these in place, the legacy **Private Money**, **AI Money**, and **Proof/Attestation Money** labels become useful analytical lenses on settlement, compute, and proof-service demand. Their Work Credits and capacity instruments remain typed service claims, not monetary instruments. Only the separate base asset may become a monetary candidate if all conditions pass. Part V climbs to **Layer 6 – Governance & Telemetry**, where neutrality and repression-resilience are kept falsifiable, not just promised.

Part V
Governance, Telemetry & Neutrality

Parts I–IV argued that the old soft guarantees are failing; that the marginal buyer of what replaces them is increasingly a delegating machine; that Privacy, Proofs, and Compute are distinct services that may support a conditional base-asset monetary candidate; that a seven-layer cypherpunk stack can supply these services; and that Layers 4–5 turn work into receipts and receipts into private, non-custodial settlement.

Part V looks at the same system from the **operations and governance** angle. If the base asset is to remain a store-of-value candidate under repression, the stack's service delivery and holder-side agency must be measured rather than promised. So must the question of who holds the asset, and on whose behalf. Work Credits remain typed service claims; their usage cannot substitute for native monetary evidence.

Chapter 25

Layer 6: Governance & Telemetry

Layer 6 is the **control plane** and the **nervous system** of the stack.

Everything below it—hardware, comms, distribution, identity, proofs, settlement—can drift, centralize, or quietly break without anyone noticing until it is too late. Layer 6 refuses to let that drift remain invisible. It insists that:

- **Neutrality** (no favored flows, no hidden house edge),
- **Repression-resilience** (still works under capital controls and censorship), and
- **Verification economics** (VERIFYPRICE, VerifyReach, VERIFYSETTLE)

  are **service-level objectives (SLOs)**, not marketing copy.

25.1 "No Dashboards, No Trust": Public SLOs as Constitution

Most protocols ship documents called "constitutions." In practice, the real constitution is whatever you cannot silently violate without getting caught.

For this stack, the constitution is:

- a **set of SLOs** on triad capacity and neutrality, plus
- a **set of dashboards and receipts** that make violations obvious.

Examples:

- "p95 VERIFYPRICE for canonical workloads stays below T seconds on commodity hardware."
- "Top-N prover/LP/router share remains below $X\%$ by hashpower/FLOPs/liquidity."
- "BTC↔XMR corridors achieve $\geq 95\%$ success with no protocol-attributable loss."
- "At least K distinct jurisdictions and hardware profiles are actively verifying."

These are not just operational goals; they are the **monetary invariants**. Layer 6's first job is to **publish and maintain these SLOs as public contracts**.

25.2 Control Surfaces: Parameters, Upgrades, Emergency Powers

Every protocol has knobs. The question that decides whether it is a constitution or a clique is not *which* knobs exist but *who can turn them, with what evidence, and how visible the turning is*. The second job of Layer 6 is to inventory the control surfaces — to name the knobs, then bind each one to a proposer, a ratifier, a data requirement, a delay, and a rollback path. A knob that is not on this list is not governed; it is somebody's private steering wheel.

The knobs come in three families. **Economic parameters** are the ones that touch money directly: block reward curves and Work Credit issuance schedules, fee policies (floor, burn, and the split between security and dividends), and PoUW weighting across workloads. **Technical parameters** shape what the network will accept: circuit versions and proof-system parameters, additions and deprecations to the accepted hardware profiles, and corridor policies such as timelocks, refund windows, and anonymity thresholds. **Operational parameters** govern the watching itself: telemetry granularity and reporting requirements, incident severity levels and response playbooks, and key ceremonies with their quorum sizes. The families differ in blast radius, which is exactly why they are enumerated separately rather than dumped into one list — issuance is not incident management, and pretending they are the same class of decision is how emergency powers leak into monetary policy.

For each surface, Layer 6 specifies: (1) who can propose a change, (2) who can ratify it, (3) what data must be presented, (4) what delays and rollback mechanisms exist, and (5) which changes are "emergency powers."

25.3 Bell-Labs-Style R&D vs. On-Chain Governance vs. Off-Chain Norms

There are three governance "forces" that need to be reconciled:

1. **Bell-Labs-style R&D.** A research org with autonomy to explore new circuits, proof systems, hardware profiles, and corridor designs.
2. **On-chain governance.** Voting weighted by governance weight (defined in Section 25.5: time-locked base asset, never Work Credits), with parameter changes baked into protocol logic. Work Credits are excluded from the franchise deliberately: they retire on redemption, so a Work-Credit electorate would be a rolling census of whoever has not yet consumed their service, and a prover could vote with credits it sold against its own capacity that morning.
3. **Off-chain norms and institutions.** Foundations, labs, dev collectives, and community norms.

Layer 6's thesis is not "pick one." It is:

> **Use Bell-Labs-style R&D to discover, on-chain governance to ratify and constrain, and off-chain norms to fill the gaps—but keep all three under telemetry.**

Concretely:

The R&D lab:

- Maintains public roadmaps and risk registers.
- Publishes upgrade proposals with quantified impacts on VERIFYPRICE, VerifyReach, VERIFYSETTLE, decentralization, and hardware profiles.
- Runs testnets and shadow deployments.

On-chain governance:

- Controls **scarce resources**: issuance, reward splits, canonical workloads, acceptance/deprecation of proof systems and hardware profiles.
- Is bounded by **constitutional SLOs**: certain changes are simply not allowed if they push metrics beyond thresholds unless higher-order safety processes trigger.

Off-chain institutions:

- Operate under **public charters** that explicitly state their mandate and constraints.
- Are expected to publish minutes, risk assessments, and incident reports.

This three-body system is inherently unstable; Layer 6 keeps it from drifting into pure plutocracy or pure priesthood by insisting that **all three bodies are visible in telemetry**:

- Changes in code and parameters appear on-chain.
- Lab work appears in open repos and benchmarks.
- Institutional decisions appear in charters, public calls, and reports.

No dark corners; no "just trust us, we're the stewards."

Governance Capture Risk

Token-weighted governance creates capture risk: large holders can centralize control over issuance, fee routing, and telemetry rules. The stack mitigates this through hard constitutional constraints that governance cannot override (see Section 25.5), timelocks on parameter changes, independent measurement infrastructure, and non-governable red lines (Section 30.6). If governance can alter issuance, fee routing, or telemetry rules without hard constraints, this is itself a red line.

25.4 Governance as SLOs: The Five Invariants

Governance and operations exist to keep one invariant true under stress:

> **Pay the machine only for work anyone can verify cheaply—and give humans lawful privacy by default.**

Everything below—roles, metrics, runbooks, and legal posture—turns that sentence into procedures, with receipts. Where policy pressure rises, the system defaults to measurement and non-custodial design rather than new gatekeepers. "No dashboards, no trust" is not a slogan; it is the rule for deciding when the system is still safe to treat as money.

Governance and operations are about what happens when the world pushes back. They are the difference between a beautiful mechanism that works in a friendly lab and an actual monetary substrate that survives YCC (yield-curve control), capital controls, app-store bans, and "secure enclave" mandates.

The protocol's "constitution" reduces to **five invariants**:

1. **Verification asymmetry, by modality.** The invariant is stated in the modality terms of Section 22.2.1, because the aggregate ratio does not protect the property the monetary claim needs. Three clauses. (a) The **M1+M2 share of verified units** — the fraction of the network's verified output whose correctness anyone can check cheaply — stays above a published floor, and that share is reported as its own series on the Value Capture Board (Section 26.9). (b) Each canonical workload's $r(W) = v(W)/p(W)$ stays within its modality band: $\leq 10^{-3}$ for M1 SKUs, $\leq 10^{-2}$ for M2 SKUs with the soundness parameter published beside the ratio; M3 SKUs carry no $r(W)$ target because the ratio is not their operative quantity. (c) The aggregate $r(W) \leq 0.3$ is retained *only* as a portfolio-level engineering floor for workloads not yet assigned a band, and no monetary claim reads on it. The **Sev-1 condition** reads on clauses (a) and (b): the M1+M2 share falls below its floor, or a banded workload's $p_{95,t}$ breaches its constitutional VERIFYPRICE SLO (Section 22.6) for the published window. A workload drifting above 0.3 is an engineering incident, not a monetary one; a single threshold such as $r \leq 0.3$ is the tempting place to put the economic hinge, and Section 22.2.1 shows why it is the wrong one.
2. **Open admission.** Anyone who brings correct work clears. Admission is not gated by identity, licensing, or proprietary hardware. Routers are neutral and open-source; house share and time-to-first-fill for new provers/LPs are visible on the Neutrality & Admission Board. If honest new entrants cannot join and get paid within a bounded window, decentralization is already drifting.
3. **Useful-work security budget.** Block rewards and fees underwrite **useful work**—MatMul-PoUW, verified inference, provenance and settlement proofs—under explicit SLAs. The security budget is tied to canonical workloads with published VERIFYPRICE rather than to a

puzzle with no external buyer. That is a deliberate trade and not a free improvement: a work function with a buyer has a demand curve somebody can move, which is the objectivity a buyerless hash puzzle gets for nothing (Section 34.1.1).

Inference Verification Tiers and WC Eligibility: Under the "AI Money" analytical lens, inference Work Credits remain service claims. To prevent weak spot checks from being marketed as high-assurance verification, workloads are tiered by verification strength:

Tier	Verification Type	WC Eligibility
Tier A	Cryptographic (full ZK)	Full eligibility ($1.0\times$)
Tier B	Probabilistic (bounded error, audited)	Discounted ($0.6\times$)
Tier C	Attestation-only (TEE + sampling)	Service-grade only ($0.3\times$); collateral-ineligible

Table 25.1: Inference verification tiers and WC eligibility.

Rule: Tier A and Tier B verified inference are eligible for full and discounted Work Credit issuance respectively; Tier C mints service-grade credits at $0.3\times$ that are ineligible as collateral (Section 22.10). Tier B issuance is discounted by a published *issuance weight* (the $0.6\times$ of the tier table in Section 22.10) reflecting error bounds and audit rate — a weight on what is minted, not a haircut on what is pledged. Tier C can exist as a service market but is not collateral-grade—it represents trust in TEE vendors, which is exactly what the thesis aims to minimize.

4. **Lawful privacy by design.** Settlement is neutral and non-custodial by default—adaptor-sig atomic swaps, shielded pools, privacy rails—but comes with **viewing keys and auditable PIDL receipts** so law-abiding users can evidence obligations without re-introducing chokepoints. The Settlement & Privacy Board reports swap success, refund safety, and anonymity-set health; flows that never touch custodians should still leave enough receipts that auditors can do their jobs.
5. **Verifiable machines at Layer 0.** Canonical proving and attestation paths run on open designs with sampled supply chains; hardware attestation becomes one input to proofs, not a vendor priesthood. Hardware profiles, lot-sampling coverage, and profile incidents show up on the Hardware/Layer-0 panes of the boards.

These invariants are **hard to change and easy to measure**. Everything else—fee curves, circuit versions, work-function parameters, corridor lists—is configuration.

To keep configuration from degenerating into a governance circus, changes are tied to **observable triggers** instead of vibes:

- If **VERIFYPRICE p95** for a workload W drifts above its target band for a sustained window, that workload's circuit and parameters are queued for revision. Blocking the change requires an explicit override that cites alternative SLOs and appears in the governance log.
- If **entry latency** for new provers/miners/LPs rises beyond a threshold, or the **top-N share**

and house share breach caps, neutral routers automatically ratchet down house share and prioritize small bidders until the metric recovers.

- If a **swap corridor** records a protocol-attributable loss — a party with no no-loss exit within τ_C (Section 23.4) — in VERIFYSETTLE, that corridor is marked *advisory-delisted* in the published route table: default clients stop routing to it, and any client may override the advisory with a visible flag. A hard removal that clients cannot override requires the breach to be reproduced in the public regression suite. Re-listing requires the suite to pass (including stress tests) plus a public post-mortem. User-liveness failures do not trigger this path; they are published under their root-cause class.

In this frame, "governance" does not micromanage every adjustment; it **writes the SLOs and the triggers**. Operations enforces them and publishes the receipts. Most changes are **data-driven roll-forward**: parameters evolve in response to Verify* drift, with clear before/after numbers.

To keep rule-making separate from rent-collection, two broad roles are distinguished:

- A **spec & telemetry steward** that publishes ABIs, PIDL schemas, canonical workloads, hardware profiles, and the VERIFYPRICE/Reach/Settle dashboards. It does not run routers or take custody. Its mandate is to keep the bill of materials and metrics honest and up to date.
- **Market participants**—miners, provers, routers, corridor LPs—who compete on price and reliability under those public SLOs, with SLA escrow and slashing keyed to PIDL receipts. Their incentives are economic; their constraints are the invariants and metrics.

This "two-chamber" structure keeps the job of defining rules and SLOs distinct from the job of selling capacity under those rules. It avoids the familiar pattern in which the entity that can edit parameters also happens to own the biggest fleet of nodes.

Governance as SLOs, not personality, is what lets the system answer policy pressure with dashboards and runbooks instead of press releases. When asked "how do you survive X?", the right answer is not "trust the foundation," but "look at this board, this trigger, and this incident playbook."

25.5 Constitutional Enforcement

"SLO-bounded governance" is only credible if the enforcement mechanism is specified.

25.5.1 Machine-Enforced Constraints (Hard)

Certain invariants are enforced **on-chain or in protocol code**:

The house-share cap is not machine-enforceable without an identity assumption. A router contract that "rejects bids exceeding share" sounds enforceable. A router can only reject a bid

Constraint	Enforcement
Base-asset issuance cap per epoch	Protocol rejects issuance exceeding the envelope of Section 25.6.1
Work Credit reservation bound	Protocol rejects a sale whose weighted units exceed committed capacity, or whose reservation would make outstanding claims on shared capacity infeasible (Section 25.6.2)
Corridor refund timeout	Atomic swap scripts enforce timeout
House-share cap	*Telemetry and slashing, not a router rejection*: see below
VERIFYPRICE bound violation	Workload suspended from WC eligibility

Table 25.2: Machine-enforced constitutional constraints.

it can attribute, and attribution to a *house* requires knowing which provers are the same party — a Sybil-resistant identity that Layer 3 deliberately does not provide for market participants. Without it, a cap on house share is a cap on the share of any single *key*, which an operator satisfies by running more keys. The cap therefore belongs to telemetry-and-slashing rather than to the hard table: house share is estimated from receipt correlation (co-location, timing, shared FCR references, shared stake sources), published on the Neutrality & Admission Board with its estimation method, and a router whose estimated house share exceeds the cap is slashed and de-prioritized by the trigger of Section 25.4. If a future Layer-3 design supplies Sybil-resistant operator identity, the cap can return to the hard table with that assumption stated in the row.

25.5.2 Override Path (Slow, Expensive)

Some constraints need flexibility. The override path is **deliberately expensive**:

- **Supermajority:** $\geq 67\%$ of *governance weight*, defined as base-asset units voluntarily time-locked for the vote's duration plus the override's timelock, weighted linearly by quantity and capped per address family at the Red Line 8 concentration threshold. Work Credits carry no governance weight (Section 25.3); base asset that is not locked carries none either, so a holder who wants to sell into the decision cannot also make it.
- **Long timelock:** ≥ 30 days (90 days for issuance-related).
- **Risk memo:** Mandatory written analysis, hash anchored on-chain.
- **Sunset clause:** Override auto-expires after 12 months.

The asymmetry is deliberate. Following the constitution costs nothing, and departing from it costs a supermajority, a month, and a memo on the permanent record, because a protocol that bends easily is a protocol someone else will bend. The price of that design is that a wrong threshold persists, and a threshold is never wrong on its own: someone chose it. Every threshold the protocol enforces is therefore published beside the reasoning that set it, and those who set a threshold own its errors as fully as the signers of an override own theirs. The override record exists to judge the rule across many decisions—whether departures under stated condi-

tions did better or worse than the rule would have—and not to convict a signer of a bad draw. A governance process that treats every override followed by a loss as a failed decision will get obedience, and obedience to a wrong threshold is the failure the telemetry was built to see.

25.5.3 Emergency Path (Narrow)

True emergencies (proof system break, active exploit) require faster action. The emergency path is **narrow and auditable**:

Allowed	Not Allowed
Deprecate proof system (hard)	Increase issuance
Delist corridor (*advisory by default*; hard only on a reproduced refund-safety breach)	Confiscate user funds
Quarantine hardware profile (hard)	Bypass refund safety
Freeze workload class (*advisory by default*; hard only on a reproduced acceptance-bound break)	Mint unbacked credits

Table 25.3: Emergency path scope limitations. "Advisory" means the action changes the published default route or workload table and every client may override it with a visible flag; "hard" means protocol code refuses the object.

Requirements: ≥ 3 of 5 security signers; auto-expire in 7 days; postmortem within 14 days.

Subtractive is not harmless: the compelled-signer problem. It is tempting to treat every subtractive power as safe because it cannot redirect value. That is true of theft and false of censorship. A body that can delist a corridor or freeze a workload class on three signatures is, to a state that can compel three signers, a kill switch for exactly the routes the stack exists to keep open — and Invariant 2 (open admission) and the "no kill switch for any single jurisdiction" commitment of Section 27.3 are both violated the first time it is used that way, with every individual action constitutional. Two of the four powers are therefore **advisory by default**: an emergency delisting or freeze changes the default table that clients ship with, is logged with its evidence, and can be overridden by any client operator who is prepared to carry the flag. The hard form — protocol code refusing the corridor or workload — is available only when the triggering breach has been *reproduced in the public regression suite* (a refund-safety loss on the corridor; an acceptance-bound break for the workload class), so that the thing being removed is demonstrably broken rather than merely disfavoured. Proof-system deprecation and hardware-profile quarantine remain hard, because their triggers (a cryptographic break, a compromised attestation root) are themselves publicly verifiable artifacts and there is no honest client that wants to keep verifying against a broken proof system.

Reconciling the emergency path with Red Line 8. Section 30.6 retires the thesis if governance can alter issuance, fee routing, or telemetry rules without hard constraints. Read care-

lessly, this path appears to grant exactly that: five signers on a seven-day fuse. The reconciliation is structural and worth stating as a constitutional rule rather than leaving to inference:

- **Emergency powers are subtractive only, and the censoring subset is advisory.** Every allowed action removes or restricts a protocol claim (deprecate, delist, quarantine, freeze). No allowed action redirects value: no fee re-routing, no issuance change, no collateral-rule change, no telemetry amendment. A body that can only shrink the protocol cannot steal from it; the worst case is a smaller protocol, which is a survivable failure and the correct bias for a seven-day fuse — provided that the two powers which shrink the protocol *selectively* (delist, freeze) cannot be made to bind on clients that disagree, which is what the advisory default secures.
- **Subtractive is not free.** Quarantining a hardware profile concentrates verification (Red Line 3's condition); freezing a workload class cuts its fee flows. Emergency use therefore carries its own falsification weight: emergency actions must be logged against the red-line conditions they aggravate, and *repeated or routine* use of subtractive powers — the "emergency governance becoming routine" failure named under Red Line 12 — trips Red Line 8's spirit even though each individual action was constitutional. The test is not whether any single action was permitted but whether the emergency path has become a de facto governance channel.
- **Signers are named, accountable, and rotating.** The five seats are filled by published charter, occupied by independent security researchers and operators (not the foundation, not the largest holders), subject to rotation, and removable by the standard override path. Anonymity of the signer set is itself a Red Line 4 telemetry failure.

25.6 Monetary Constitution

It is natural to write that "Work Credit issuance is governed by a halving schedule." That sentence conflates the two instruments the whole document works to keep apart. A halving schedule is a property of a *base asset*: a fixed, front-loaded supply path that no demand can move. A Work Credit is a typed claim on reserved future service; it is issued when committed capacity is sold and reserved and extinguished when the service is delivered, and there is no sense in which the number of such claims can "halve" on a calendar without either refusing to sell committed capacity or selling claims on capacity nobody committed. The two have separate subsections with separate rules.

25.6.1 Base-Asset Issuance Envelope

Gross base-asset issuance follows the consensus schedule fixed ex ante — halving every N blocks (e.g., 4-year halvings). This is the supply path of the monetary candidate and noth-

ing else. Audited Delivered Verified Capacity neither expands nor contracts it. Net supply additionally reflects the retirement rules of the fee split (Section 25.6), so net supply is not predetermined while burns depend on usage. Capacity evidence constrains service claims (Section 25.6.2), not the monetary issuance envelope.

Why there is no capacity modulator. A rule that moved issuance with audited capacity growth would remove one kind of discretion and import a dependence in its place: issuance would rest on measurement choices, audit judgments, workload aggregation, scenario definitions, and a target growth rate — the measurement problem of Section 17.9, carried into the monetary object. It would also not answer the question that matters for the base asset, which is whether a held balance outperforms Bitcoin-funded procurement (Section 13.1). The reference design is simpler because the mechanism has not earned its place, not because a fixed schedule proves monetary value. Removing it settles neither the adequacy of the security budget, nor the incentives of useful-work consensus, nor governance over the other parameters.

25.6.2 Work Credit Issuance Rules

Work Credits have no schedule. They are issued under four rules and no others:

1. **One claim per standardized unit of committed capacity.** A Work Credit is issued only against a registry SKU (Section 22.3.3) committed by a named provider or constituted pool for a stated window, and the sale and the reservation are one atomic step: no sold claim exists without its capacity reserved. Verified delivery against the SKU's acceptance criterion extinguishes the claim and issues the Work Receipt. There is no issuance "to" anyone; there is only sale against reserved capacity.
2. **Weights are at most one.** The issuance weights of Section 25.6.5 (verification tier, hardware grade, concentration, FCR confidence, sovereign optionality) multiply the issued quantity and are all ≤ 1.0. No grade, tier, or bonus issues more credit than capacity committed; premia for open hardware or high assurance are paid through collateral grade and routing priority, never through issuance above committed capacity.
3. **Reservation-bounded.** The weighted Work Credits outstanding against an SKU class cannot exceed the reservation limit set by audited stress-adjusted Delivered Verified Capacity, computed jointly across workloads that share capacity; a sale that would breach the bound is rejected by protocol code (Section 25.5). DVC supports this limit; it does not guarantee performance.
4. **Failure resolves by terms.** Cancellation, expiry, default, and refund follow the terms written in the claim. A failed delivery becomes neither an open-ended promise nor a claim on the base asset, and advance proceeds that finance capacity not yet built are project finance, labelled as such (Section 3.3).

   Nothing in these rules references the base asset's schedule, and nothing in the base asset's

envelope references Work Credit volume or capacity measurement.

25.6.3 Fee Routing

Destination	Share	Mechanism
Capacity providers and their delegating stakers	70%	Paid via SLA escrow; delegation terms published
Protocol burn	20%	Permanently removed
Assurance budget	10%	Funds Layer 0 sampling (the Assurance Fund of Section 17.5.4), audits, security review

Table 25.4: Fee routing split — the reference-design partition. Every worked example in Parts II–V uses this point; the 30–50% burn band of Section 8.4 is the design space it was chosen from, and Design A's 40/30/30 is an alternative inside it, not a second reference.

Why burn, and what a burn is. A burn returns fee value to holders pro rata: it is a buyback, and Section 13.3.9 and Section 27.2 both say so. The familiar chain "high demand → high fees → high burn → supply reduction → value appreciation; this closes the loop" is the inference Section 13.3.9 shows to be false, and Layer 6 does not make it. What a burn does, exactly: it reduces *net* issuance — and reduces supply outright only in an epoch where burned units exceed issued units, which at the Year-1 fee-coverage target of Section 25.6.6 they will not (if retained fees cover 30% of the security budget, gross fees are 37.5% of it and the 20% burn share is 7.5% of the budget, against issuance funding the other 70%). It is published as its own series — gross burns, gross issuance, net issuance — on the Native Monetary Buyer Map, and it is read as supply arithmetic, not as buyer evidence (Appendix H.14). It does not by itself create monetary premium: a buyback accrues value to a claim, and accrual is the cash-flow leg of Chapter 8, not the moneyness leg.

This routing is what makes the Value Capture Lemma from Chapter 8 more than a slogan: it is the fee plumbing that delivers the lemma's cash-flow claim, and only that claim. Its 20% burn is the same conservative variant used in the Layer 4 worked example (Chapter 22) rather than the 30–50% reference band of Chapter 8, because here the remaining 10% funds assurance. The assurance budget funds the Layer 0 sampling that keeps hardware claims credible. Changes to the fee split are governance parameters — and every such parameter is discretion over the return to holders that the classification analysis examines (Section 27.2); the constitutional version of this design fixes the split *ex ante* and removes the knob rather than defending a floor for it.

25.6.4 Retirement Rule

In addition to fee burn, base-asset units are retired in specific circumstances:

Trigger	Retirement
Work Credit redemption	Any base-asset fee paid to fulfill the capacity claim is retired under the fee rule; the Work Credit itself is retired as a fulfilled service claim
Slashing	100% of slashed collateral is burned (not redistributed)
Failed workloads	Fees for failed proofs (prover fault) are burned, not paid

Table 25.5: Base-asset and service-claim retirement triggers.

25.6.5 Issuance Weights

Work Credits minted under weaker assurance levels receive discounted issuance. The multipliers below are *issuance weights* — factors on what is minted — and are named as such; the word *haircut* is reserved in this document for discounts on what is pledged (collateral grade, bridge class), which are a different instrument applied at a different moment. Calling this table "Risk Haircuts" would not be a cosmetic error: a haircut leaves the claim whole and discounts its pledge value, while a weight mints fewer claims.

Factor	Issuance weight	Example
Hardware profile (L0 grade)	L0-A: 0.9×; L0-B/C/D: 1.0×	Lower-assurance hardware earns fewer credits; no grade earns more than one credit per unit committed (Section 17.4)
Verification tier	Tier A: 1.0×; Tier B: 0.6×; Tier C: 0.3×	Probabilistic verification earns less than cryptographic
Prover concentration	>20% share: 0.9×; >30%: 0.8×	Concentrated provers earn progressively less
FCR confidence grade	Missing/unverifiable FCR: issuance cap; low-confidence FCR: 0.85×	Facilities that cannot substantiate physical capacity claims earn fewer, capped credits
Sovereign optionality ($\mathcal{O}_s$)	High band: 1.0×; moderate: 0.9×; low: 0.75× and issuance cap	Facilities whose physical capacity is fully substantiated but structurally fragile earn fewer credits

Table 25.6: Issuance weights for Work Credit minting. Every weight is ≤ 1.0 (Section 25.6.2).

These weights are protocol parameters (not discretionary); they appear in dashboards and apply automatically.

Two distinct physical factors. The last two rows look similar and do different work. The distinction matters enough to state explicitly, because collapsing them leaves a gap that a well-run fragile operator can walk through.

- **FCR confidence grade is epistemic.** It asks whether the physical substrate can be seen at all. It penalizes opacity.
- **Sovereign optionality is strategic.** It asks whether what can be seen is survivable. It penalizes fragility (Section 17.8).

Without the second row, a facility that reports complete, signed, independently verified FCR data documenting a single interconnect, one fuel source, no backup duration, and a jurisdiction with an active rationing regime would receive *no issuance discount whatsoever*. It would in fact score well, because its disclosures are impeccable. Transparency about fragility is not the same as resilience, and an issuance schedule that rewards only the former is measuring candour rather than durability.

This is the operational form of the distinction drawn in Section 17.10.1: verifiability and availability are separate properties, and Work Credit issuance should price both.

Interaction with the affordability hinge. These weights must not become a back door for degrading verification cost. Per Section 17.10.3, a resilience premium is acceptable only while benign-state Real-Resource VerifyPrice stays inside its constitutional SLO band. If pursuing a higher $\mathcal{O}_s$ band pushes a network's verification cost outside that band, the SLO wins and the optionality target is deferred. Weights adjust issuance; they do not license breaching a constitutional target.

25.6.6 Coverage Target

The **fee-coverage ratio** measures what share of the security budget comes from fees vs. issuance:

$$\text{FeeCoverage} = \frac{\text{RetainedFees}}{\text{TotalSecurityBudget}}$$

where RetainedFees is fee revenue net of the burned share — the portion that actually funds validators, provers, and operations. Burns do not fund security: a burn is supply arithmetic, not security revenue (Appendix H.14), so adding it to the numerator would double-count the burned slice of fees (it is already inside gross fee revenue) or, if read as destruction of issuance, would count destruction as funding. The burned quantity is published as its own series and read as reduced net issuance. Gross fee revenue is likewise published, so the burn share is visible; the convention is pinned here and applies wherever a fee-plus-burn quantity appears (the take-rate τ of Section 30.6 uses *gross* fees, because it measures what buyers pay, not what the budget retains).

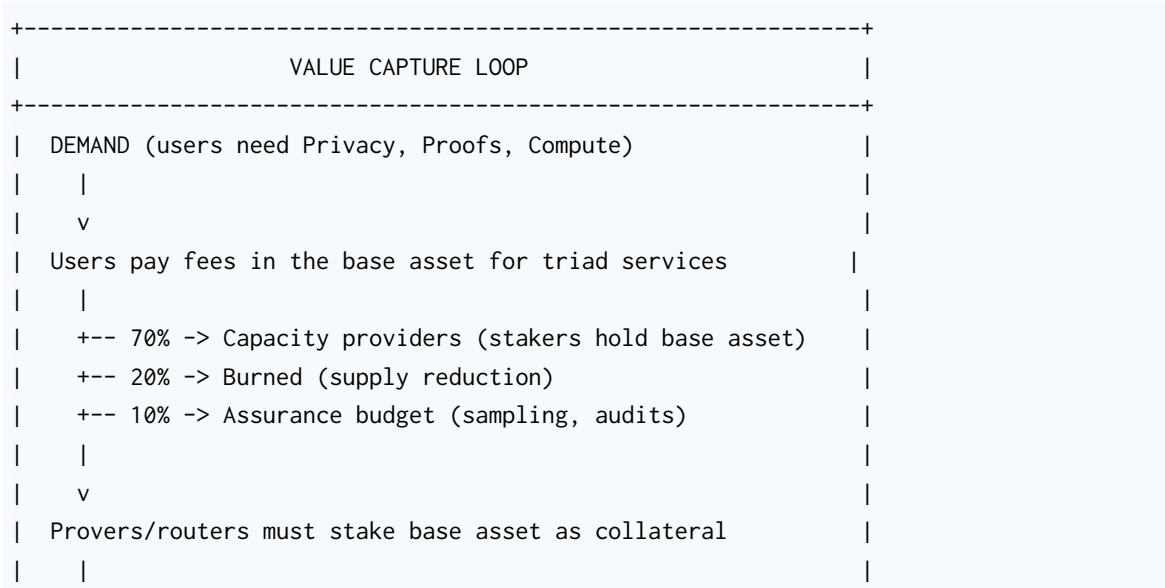
Why this matters: A system funded entirely by issuance is inflating its way to security — the warehouse paying guards in shares printed that morning (Section 30.2 gives the analogy its full form). A system where fees cover the budget has structural demand: someone is paying for the service, not merely holding the token. The trajectory from issuance-funded to fee-funded is the path from "speculative token" to "productive asset," and it is a trajectory, not a switch, precisely because trust in the funding model has to be earned quarter by quarter.

Target	Timeline
FeeCoverage $\geq 30\%$	Year 1
FeeCoverage $\geq 50\%$	Year 3
FeeCoverage $\geq 80\%$	Year 5+

Table 25.7: Fee coverage trajectory. The staged targets apply *conditional on scale*: they bind from the first quarter in which native-loop volume exceeds the published minimum — 1,000 daily active settlement identities and USD 5M quarterly native-settlement volume, the same viability gate the boards use — so the Year-1 row means "within four quarters of viability," not "four quarters from genesis," when bootstrapping runs longer. If the gate has not been met, FeeCoverage is published with the staging flag "pre-viability" and no target is credited or missed.

The targets are a hypothesis, and the comparators say an aggressive one. No precedent supports 30% in year one. Bitcoin, after seventeen years, funds its security budget from fees in the low single digits per cent in a typical epoch, with brief spikes into double digits during fee events. The utility-token cohort the thesis engages in Section 34.1.5 — Filecoin, Render, Akash [Messari 2024–26] [Token Terminal 2024–26] — reached low-single-digit fee coverage against issuance at maturity, not 30%, and none has approached 50%. The staged table is therefore labelled for what it is: a falsifiable hypothesis that a work function *with a buyer* can reach fee coverage an order of magnitude above every buyerless and every utility-token precedent, within four quarters of viability. If the Year-1 row is missed, the miss is published against those comparators and the phase does not advance; the table is not revised downward to meet the data.

25.6.7 Value Capture Loop



```
|    v                                                              |
|  Base-asset gross issuance follows the ex-ante schedule           |
|  (Work Credits issued separately, one per reserved unit)          |
|    |                                                              |
|    v                                                              |
|  NET EFFECT: Demand -> Fees -> Burns + Staking -> Accrual         |
|  Monetary premium still requires holder and agency conditions     |
+-------------------------------------------------------------------+
```

Why demand does not simply expand base-asset supply proportionally:

1. **Issuance cap is hard:** Protocol enforces the epoch-level envelope of Section 25.6.1 regardless of demand; Work Credit volume does expand with demand, by design, and is the other instrument.
2. **Burn is automatic:** Higher demand → higher fees → higher burn → lower *net* issuance. Not lower supply until burn exceeds issuance, which the coverage trajectory does not reach before Year 5 at the earliest; and a buyback, not a moneyness mechanism.
3. **Staking locks supply:** More provers/routers → more collateral locked → less circulating supply.
4. **Issuance weights constrain minting:** Weak verification or concentrated provers sell fewer Work Credits per unit committed, so service-claim volume grows more slowly than raw capacity; this constrains the claim instrument, not the base asset.

This is the "bond indenture" that makes the accrual claim auditable, not asserted. It is not, by itself, the SoV claim; that requires the holder-side and agency conditions the diagram's last line names.

Chapter 26

Extended Telemetry

A constitution nobody can check is a press release. Layer 6 claims that measurement is the constitution — that SLOs, not assurances, decide whether the asset is still money — and that claim is only as strong as the measurement regime beneath it. Dashboards can be gamed, and the gaming is usually boring: a favorable region here, a fast machine there, a bad month quietly absent. So the observatory must be built the way it audits others. Anyone can re-run the audit, on hardware they own, and catch the observatory lying. That is the standard this chapter engineers for, and everything in it — reference machines, cost vectors, sampling seeds, anchored archives — exists to make that re-run possible.

26.1 Reference Verifier Classes

"Anyone can verify" is a promise about hardware, so it has to be priced on hardware a person can actually own. All Verify* metrics are therefore measured on **standardized reference hardware**:

Class	Hardware Spec	Use Case
Laptop-Class	4-core CPU, 16GB RAM, SSD, no GPU	Default for SLOs
Mobile-Class	ARM SoC, 8GB RAM	Edge/mobile
Server-Class	32-core CPU, 128GB RAM	High-throughput

Table 26.1: Reference verifier classes.

Rule: All constitutional SLOs (e.g., "VERIFYPRICE $t_{95} \leq 5$s" — the p95 *time* field of the canonical tuple, Appendix A) are defined against **Laptop-Class** unless a workload's registry entry binds it to another class, in which case the entry names the class. This ensures the "anyone can verify" promise is falsifiable. Where a registry entry binds an SKU to Server-Class (e.g. `INFER_LM_7B_512TOK`, whose M3 re-execution does not fit a Laptop-Class SLO; Table 22.4), the entry governs that SKU and the Laptop-Class figure is still published beside it; the blanket rule governs everything the registry does not rebind.

26.2 Reference Verifier Implementations

Requirement	Specification
Independent implementations	≥ 2 independent verifier codebases per canonical workload tier
Open source	All reference verifiers must be open-source with reproducible builds
Versioned and hashed	Binary hashes are part of the canonical workload registry
Deterministic	Same proof + same verifier version $\rightarrow$ same result, always

Table 26.2: Reference verifier implementation requirements.

Why two implementations rather than one trusted one? Because a single verifier codebase is a single point of failure for the truth itself. One bug can inflate or deflate every measurement the constitution depends on, and nobody outside would ever know. Two independent implementations, built by different people from the same specification, turn that silent failure into a visible disagreement — and a disagreement, unlike a bug, shows up on a dashboard.

26.3 Cost Vector Definition

A single number is easy to quote and easier to flatter. VERIFYPRICE is deliberately not one, and it is also deliberately not *several*: it has one definition, the canonical five-field tuple $(p_{50,t}, p_{95,t}, p_{50,c}, p_{95,c}, \text{fail})$ of Appendix A. What this section defines is the **cost-vector view** — the per-run physical quantities the harness records for workload W on verifier class V, from which the canonical tuple is priced by the rule of Section 22.6:

$$\text{CostVector}(W, V)_{\text{run}} = (t_{\text{run}},\ e_{\text{run}},\ m_{\text{run}},\ b_{\text{run}},\ \text{outcome}_{\text{run}})$$

Where:

- t_{run}: verification **time** (seconds); percentiled into $p_{50,t}$, $p_{95,t}$
- e_{run}: **energy** (Joules); enters the per-run USD cost c_{run}, which is percentiled into $p_{50,c}$, $p_{95,c}$
- m_{run}: peak **memory** (MB) — *published, not priced*; exceeding the registry ceiling is a failure outcome
- b_{run}: **bandwidth** (bytes) — *published, not priced*, same rule
- $\text{outcome}_{\text{run}}$: pass or failure code; aggregated into fail

This vector is not written as "VerifyPrice$(W, V) =$", because it is not VerifyPrice. Only the canonical five-field tuple of Section 22.2 carries the name; this vector and the per-run price model of Section 22.6 are views of its inputs, and they are labelled as such wherever they appear.

26.4 Sampling and Anti-Cherry-Pick

The most common attack on telemetry is not fabrication; it is curation. Nobody has to invent numbers when selective reporting will do: show receipts from good regions, fast hardware, favorable conditions, and the dashboard flatters itself without a single lie in it. The sampling regime exists to make curation detectable:

Mechanism	Specification
Public randomness beacon	Receipt selection seeded by on-chain randomness (e.g., block hash, VRF output)
Stratified sampling	Samples drawn proportionally from: workload class, hardware profile, geographic region, time-of-day
Adversarial corpus	10% of test proofs are malformed or worst-case (max witness size, pathological inputs)
Continuous measurement	Rolling 24-hour windows; no periodic snapshots that can be gamed
Multi-operator	≥ 3 independent measurement operators; divergence $> 5\%$ triggers investigation

Table 26.3: Sampling mechanisms for Verify* telemetry.

Rule: Any dashboard point published without a verifiable sampling seed and stratification breakdown is not part of the official Verify* record.

26.5 Data Availability and Anti-Memory-Hole

Dashboards are views. The source of truth is the **receipt corpus**. Receipts must be:

Requirement	Specification
Content-addressed	Receipt sets identified by Merkle root
Anchored	Merkle roots periodically committed to a neutral ledger (e.g., Bitcoin, Ethereum)
Archived	Raw receipt data retained for ≥ 2 years by ≥ 3 independent archivists
Queryable	Any observer can request receipts corresponding to a dashboard point

Table 26.4: Data availability requirements.

A dashboard can be edited; a Merkle root committed to a neutral ledger cannot be uncommitted. That asymmetry is the whole defense. If a government or a CDN can quietly delete the bad months, the telemetry regime has been captured without a single headline. Content-addressing plus neutral anchoring does not prevent deletion — it makes deletion *visible*, and visible deletion is itself a datum.

26.6 Reproducibility Contract

A valid dashboard datapoint is defined as:

$$Datapoint = (d_{\text{root}}, v_{\text{harness}}, v_{\text{verifier}}, r)$$

Validity check: Anyone can:

1. Fetch receipts by d_{root} (dataset root)
2. Run v_{harness} (harness version) with v_{verifier} (verifier version)
3. Verify r (result) matches

If the results do not match, the datapoint is invalid and the discrepancy is a Sev-1 incident.

Published artifacts:

- Harness code (open source, tagged)
- Verifier binaries (reproducible builds, hashed)
- Dataset roots (timestamped, anchored)
- Aggregation scripts (deterministic)

This transforms "no dashboards, no trust" from philosophy into a **reproducible measurement contract**.

26.7 How to Measure: Receipts, Not Vibes

VERIFYPRICE was introduced at Layer 4; VerifyReach at Layer 1; VERIFYSETTLE at Layer 5. Layer 6 treats them as a **coherent observability regime** rather than three unrelated graphs.

The observability story only works if it is itself verifiable. That is why every claim, proof, settlement, and SLA outcome is turned into a **receipt** with enough structure that anyone can replay the metrics on a reference verifier.

In practice, this means:

- Every proof receipt carries not only the proof artifact but also start/end timestamps, resource usage, and a hardware profile tag.
- Settlement receipts encode corridor, timing, refund status, and failure codes.
- Routers emit anonymized samples of order books, matched and unmatched bids, and house vs. third-party flags.
- Hardware vendors publish lot-attestation artifacts that receipts can reference as part of their profile.

A public "VerifyPrice Observatory" then ingests these receipts and emits the aggregate vectors referred to in the text. Anyone can pull the raw receipts corresponding to a dashboard point, feed them into the same open-source verifier suite, and see whether they obtain the same numbers. If they do not, the mismatch is itself an incident.

This architecture enforces a simple discipline: **if it matters, there is a receipt for it.** If a claim about performance, neutrality, or safety cannot be tied back to receipts that anyone can re-check, it is out of scope for the monetary thesis.

They are the three primary dials for answering:

- *Is verification still cheap?* (VERIFYPRICE)
- *Can users still reach the network under pressure?* (VerifyReach)
- *Can users still move value privately and safely?* (VERIFYSETTLE)

This observability regime watches the stack across five planes:

Proof & compute plane Are receipts reliable, and does verification asymmetry actually hold under load?

Settlement & privacy plane Do non-custodial payouts actually clear, with privacy and refund-safety matching the promise?

Neutrality & admission plane Can new provers, miners, routers, and mirrors join on equal footing, or has the system ossified into a club?

Economic coverage plane Are real fee flows, not issuance, paying a growing share of the security and operations budget?

Layer-0 hardware plane Do "verifiable machines" remain falsifiable in practice, or has the system slid quietly back into "trust the vendor"?

Each plane maps directly to the nine SoV requirements of Part I (Chapter 4): credible scarcity, cheap public verification, censorship-resistance and portability, neutrality and permissionlessness, native demand, lawful privacy by design, duration-neutrality, asset-level value capture (non-bypassability), and agency preservation. The economic coverage plane reads on the fifth and eighth; the settlement and privacy plane on the sixth and ninth; the proof and compute plane on the second; the neutrality plane on the third and fourth; the Layer-0 plane on the second and, through sovereign optionality, the third. If any plane drifts too far from its targets, the monetary thesis weakens, no matter how elegant VERIFYPRICE looks in isolation.

26.8 VerifyFlow: Instrumenting the Market Realization Plane

VERIFYPRICE, VerifyReach, and VERIFYSETTLE answer whether the stack is real. None of them answers a question that will dominate every public conversation about the asset: *why is the price doing that?* Corollary 13.1 says monetary adoption cannot be inferred from price, and Section 13.7 says the answer lives in external market structure. VerifyFlow is the fourth named verification family, and it measures the external financial representation and price-transmission state of the native monetary object.

The governing principle extends the existing one:

> No protocol dashboards, no monetary trust. No flow ledger, no price inference.

26.8.1 Wrapper-Level Vector

For each external financial wrapper i, publish

$$\text{VerifyFlow}_i = \big(L_i,\ A_i,\ \Delta A_i,\ r_i,\ w_{Xi},\ \varepsilon_i,\ \kappa_i,\ C_i,\ H_i,\ D_i\big)$$

where L_i is the leverage multiple and reset rule; A_i is assets or delta-adjusted exposure; ΔA_i is creations and redemptions, published *raw and flagged*: ΔA_i is the input to the decomposition of Appendix H.1, not an aggregate in its own right, and no downstream sum may consume it unless the wrapper's series has been decomposed into its ε-, α-, and residual components — otherwise aggregating across wrappers silently rebuilds the double-count the asset-level vector exists to prevent. r_i is the wrapper's underlying return over the daily-reset horizon, and w_{Xi} its delta-equivalent exposure weight to the native asset X (unity for a single-asset physical wrapper, less for baskets and partially hedged structures); both are consumed by the Q^{RC} and Q^{RD} terms below. ε_i is holder flow elasticity; κ_i is net mechanical gain; C_i is constituent concentration; H_i is hedge implementation (physical, swaps, futures, options); and D_i is dealer or counterparty concentration. Appendix H derives ε and κ and states their measurement contract.

26.8.2 Asset-Level Vector

For the native asset X, publish

$$\text{VerifyFlow}(X) = \big(Q^{RC},\ Q^{RD},\ \text{MPR}^{\text{gross/net}},\ \text{CCR},\ \text{SER},\ \text{WNG},\ \text{RCR},\ \text{NUS}\big)$$

The MPR slot is a published pair (gross and net), so the vector names the pair rather than one form; every other slot is a single series.

Return-coupled demand Q^{RC}. Exposure that moves as a function of the return itself, $Q^{RC} = \sum_i w_{Xi}\,\kappa_i A_i r_i$, where w_{Xi} is the wrapper's delta-equivalent exposure weight to X — the same weight Q^{RD} carries, and without it a basket wrapper's rebalancing demand for its other constituents would be booked against X. This term shapes volatility, momentum, persistence, and reversal severity.

Return-decoupled demand Q^{RD}. Allocation flow largely independent of the day's return, $Q^{RD} = \sum_k w_{Xk}\alpha_k A_k$, where α_k is the return-independent intercept of net creation into wrapper k and w_{Xk} its delta-equivalent exposure weight to X. Q^{RD} consumes α only. Per the decomposition rule of Appendix H, creations split into an ε-driven component (return-coupled, already inside κ), the intercept α_k (structural bid — distribution growth, mandate adoption, scheduled

contributions), and a residual u_k (unexplained, published as its own series and consumed by no aggregate). Feeding raw F_k into this sum would double-count the ε-component, which κ already carries through Q^{RC}, and would inflate MPR and every downstream reading built on it. This term shapes destination, concentration, and level — which prices get accepted without valuation-sensitive selling.

Mechanical Pressure Ratio (MPR). Gross measures *pressure*; net measures *direction*. That distinction is the whole metric, and it is worth the two definitions it costs. Q^{RC} and Q^{RD} partition product-level demand by return-dependence: the ε-component of creations belongs to Q^{RC} (inside κ) and only α_k belongs to Q^{RD}. Total product-level *exposure demand* is then $Q^{\text{exp}} = Q^{RC} + Q^{RD}$ with no term counted twice. Dealers translate that exposure into market orders, so realized market demand is its intermediated image,

$$Q^{\text{mkt}} = \mathcal{H}\big(Q^{\text{exp}},\ \text{dealer inventory},\ \text{internalization},\ \text{swap/future/option mix}\big).$$

Left as a symbol, $\mathcal{H}$ would make MPR a ratio with an undefined numerator. It is defined here as the published **pass-through decomposition**: for each reporting segment j (venue $\times$ wrapper class),

$$Q_j^{\text{mkt}} = \big(1 - \iota_j\big)\,\pi_j\,Q_j^{\text{exp}} \;-\; \Delta I_j,$$

where ι_j is the dealer internalization share (exposure netted against other client flow and never sent to market), π_j the spot pass-through of the segment's hedge mix (unity for physical holdings; the delta of the futures, swap, or option book otherwise), and ΔI_j the change in dealer inventory over the window. Each of ι_j, π_j, and ΔI_j is a disclosed series under Section 26.8.5; where a segment's dealers do not disclose them, Q_j^{mkt} is *not computable*, MPR is published with that segment marked missing and its exposure share stated, and no headline MPR is quoted for a window in which the missing share exceeds 20% of Q^{exp}. MPR is thus computable exactly to the extent dealers disclose, and the metric says so rather than filling the gap with an assumption.

The gross form, $\text{MPR}^{\text{gross}} = \sum |Q_j^{\text{mkt}}|/V_X$, counts every dollar of liquidity consumed on both sides: it asks how hard the machinery is leaning on the market, whichever way it leans. The net form, $\text{MPR}^{\text{net}} = |Q^{\text{mkt}}|/V_X$, nets the offsetting flows first and asks only where the weight ends up. Here V_X is the asset's executable liquidity — average daily traded USD volume on the reference venue set over the same window the numerator aggregates, published with the ratio — and the segments j are the pre-declared reporting partition (venue $\times$ wrapper class), also published with the ratio.

The two forms must never be averaged, ranked, or otherwise combined, because each guards against a different deception. Netting offsetting flows before the modulus erases the liquidity both sides consumed, which is why the gross form leads and the net form contextualizes. And customer exposure and dealer hedging are the same demand observed at two stages, not two

independent demands; summing them would double-count. Report Q^{exp} and Q^{mkt} separately and never add them — a pressure gauge and a compass are two instruments, and gluing them into one dial breaks both.

Custodial Control Ratio (CCR). Base asset controlled by custodians and wrappers, divided by circulating supply.

Synthetic Exposure Ratio (SER). Delta-adjusted derivative exposure divided by liquid free float.

Wrapper–Native Growth Gap (WNG). $\text{WNG} = \Delta \ln E^{\text{wrapper}} - \Delta \ln U^{\text{native}}$, where U^{native} counts native service *flows* only — fee payments, private settlement volume, proof purchases, and verified-compute consumption, each a turnover measured over the window. Collateral lockup is a stock and is published as its own companion series, never summed into U^{native}; burns are supply arithmetic, not buyer evidence (Appendix H.14), and are likewise excluded. A persistently positive WNG means financial exposure is growing faster than monetary use — the single most compact measure of financialization outrunning adoption. Because it feeds the triple-divergence composite and Red Line 9, its aggregation is specified rather than left to the operator: a trailing four-quarter window; quantities in *unit* terms on both sides (wrapper exposure in units of X equivalent, native use in units of service transacted), so that asset-price revaluation is stripped from the comparison by construction; and the USD-denominated series published alongside for readers who want the market view. The reason for unit-denominating is a revaluation artifact, and it runs the other way from what intuition suggests: E^{wrapper} is a stock observed instantaneously and revalued with every tick, while U^{native} accumulates at transaction prices over the window, so a price *rise* — not a fall — can lift measured $\Delta \ln E^{\text{wrapper}}$ on revaluation that no usage produced. Unit terms cancel that term; the USD series exposes it. The deflation convention here is therefore *unit denomination*, which is distinct from the *workload unit-price deflation* of Red Line 14 Condition B (a compute-price index, not the asset price): same word, different deflators, both published.

Recursive Claim Ratio (RCR). Exposure held in wrappers whose underlying is itself another wrapper or derivative, divided by total wrapper exposure — delta-adjusted on both sides, over the same window as SER. One minus RCR is the share of wrapper exposure that bottoms out in the native asset directly.

Native Use Share (NUS). The share of *observable economic activity touching the asset* that actually settles through the protocol-native monetary loop, where the universe is defined as all activity on public venues and wrappers reporting the asset (USD value of transfers and wrapper flows over the window), and the numerator is the USD value of native-loop settlement (fee-paying, proof-carrying, or corridor) over the same window. The universe is pre-declared with the metric; activity invisible to public data is excluded from both numerator and denominator and acknowledged as a coverage limit, not silently assumed zero.

26.8.3 Holder Cohorts, Vintage, and Policy Concentration

Flow elasticity is not a permanent product characteristic, and treating it as one is a modeling error. Product launch date predicts holder behavior [Green 2026a]: older funds contain embedded winners who are more likely to skim gains after a rally, whereas newly launched products may contain momentum-oriented holders with no gains to realize and therefore nothing to trim. So

$$\varepsilon_{i,t} = f(\text{product age},\ \text{embedded gains},\ \text{holder cohort},\ \text{narrative regime},\ \text{volatility})$$

Telemetry should therefore track product age, average holder cost basis where inferable, the realized/unrealized gain distribution, rolling elasticity, statistical break dates, holder-cohort turnover, and fund-flow response after drawdowns.

This matters more, not less, in an agent economy. Automated treasury systems, robo-advisers, and AI allocation agents may share common rebalancing policies, producing a machine-readable version of holder elasticity at scale. The protocol should therefore monitor concentration not only of holders but of **holder operating rules**.

Policy Concentration Ratio

The share of net asset demand governed by the largest common rebalancing templates, treasury algorithms, robo-advisers, or agent policies. Picture a thousand independent drivers, each in their own car, all running the same navigation app: ownership is distributed, but on the day the app says "take the exit," they all take the exit. An asset can have decentralized ownership and highly centralized behavior; only the second predicts what happens on a bad day.

26.8.4 Capital Survival: AUM Is Not Investor Return

One further metric prevents a specific and common confusion. A leveraged wrapper can survive commercially while repeatedly destroying investor capital: reverse splits preserve a presentable share price, fresh inflows replenish assets, and the sponsor continues collecting fees on a product whose successive holder cohorts were harmed. Persistent AUM is therefore not evidence of successful capital preservation.

Capital Survival Ratio (CSR)

At the product level, every term must be an aggregate dollar amount — *net assets*, not per-share NAV, which would mix a per-unit price into a sum of dollar flows:

$$\mathrm{CSR}_i = \frac{\text{current net assets} + \text{cumulative distributions} + \text{cumulative redemptions}}{\text{cumulative subscriptions}}$$

Because the stated purpose is to ask whether *successive cohorts* survived, the product-level figure is only a summary. The cohort form is the one that answers the question, for cohort c with cohort-attributable value $V_{c,t}$:

$$\mathrm{CSR}_{c,t} = \frac{V_{c,t} + \sum_{\tau \le t} \text{cash returned}_{c,\tau}}{\sum_{\tau \le t} \text{capital contributed}_{c,\tau}}$$

This is a multiple on invested capital by vintage; a dollar-weighted return over the same cohort carries the same information (a cohort's CSR exceeding 1 after T quarters corresponds to a positive dollar-weighted IRR over those flows, so the two are orderings of one object, not independent evidence). A product can show CSR_i near or above 1 while late cohorts sit well below 1, which is precisely the failure the metric exists to expose. Accounting implementations vary, but the purpose does not: AUM is not investor return, product survival is not holder survival, and a wrapper can be an excellent business for its sponsor and a disaster for everyone who owned it.

26.8.5 Dealer Balance Sheets and Hedge Paths

Wrapper exposure does not necessarily arrive as direct spot purchases. Dealers may implement it with swaps, futures, baskets, or options, and counterparty capacity affects how — and whether — the exposure reaches the underlying market. Product-level exposure demand must therefore never be read as an identical quantity of spot orders.

Publish: the wrapper counterparty list; top-N dealer share; physical holdings versus swaps; swaps versus options; average derivative maturity; delta- and gamma-adjusted exposure; collateral and margin requirements; evidence of counterparty caps; and changes in hedge-instrument mix.

Dealer capacity warning. A migration from swaps toward options can indicate that dealer balance-sheet capacity is binding. When dealers are short calls, rising spot and rising volatility can require additional underlying purchases, adding a gamma feedback loop on top of the mechanical one. This is an empirical warning to watch for, not a universal law to assume.

26.8.6 Convenience-Yield Telemetry: Pricing the Holder-Side Service Flow

VERIFYPRICE, VerifyReach, and VERIFYSETTLE measure the stack; the wrapper vector measures the realization plane; the Native Monetary Buyer Map measures who is holding and how. None of them prices the thing the monetary claim actually rests on. Section 13.3.11 relocates the premium to a **regime-contingent convenience yield** — a holder-side service flow with no cash flows to discount — and then concedes it is unsized. The concession is honest

but incomplete: unsized does not mean unmeasurable, and an unsized quantity with no instrument is indistinguishable from an article of faith. Storage theory [Kaldor 1939] [Working 1949] extracts convenience yield from observable prices rather than asserting it (the forward-curve slope net of carry); the crypto-native analog is the lending and derivatives complex around the asset. The thesis therefore publishes a fifth measurement family whose job is to make the premium a number rather than a hope.

Instruments. For each venue publishing a rate, report:

Base-asset lending rates. The market-clearing rate for borrowing the base asset itself, segmented by venue, term, and counterparty class. A borrower pays this rate to be short the asset; what they are paying *for* is the units. When holders of the bearer service lend out their units, they are selling the service flow for the term of the loan — so the lending rate is the direct market price of temporary access to bearability, the closest available analog to a lease rate.

Forward and perpetual basis. Forward price minus spot, and perpetual funding rate, net of the reference interest rate. Persistent positive basis that exceeds carry costs indicates demand for *deferred or levered exposure* outrunning demand to hold the units — demand for the price story rather than the bearer service. Persistent negative basis indicates the reverse: entities that need the units (operators meeting collateral calls, shorts delivering into settlement) bidding against the float.

Wrapper basis. The premium or discount of a custody-restricted wrapper over the value of its native holdings. A persistent premium is the price of *not* holding natively: legal wrapper, custody, and access convenience purchased in place of self-custody, agency, and exit. It is the complement of the lending rate — both price the friction between native holding and representative holding — and under Corollary 13.1 a large and rising wrapper premium with stagnant native use is financialization, not monetary adoption.

What the instruments read. Convenience yield is a *residual claim* on holding, extracted from prices that also contain discount-rate and risk information, so each reading requires the same care as ε in Appendix H:

- **Lending rate decomposition.** Borrow demand has three sources: directional short interest, market-making inventory financing, and operational short squeeze against locked float. Segment by borrower class where observable; an operational squeeze (collateral calls meeting immobile supply) is a liquidity event, not a premium reading.
- **Regime conditioning.** The thesis's claim is specifically *regime-contingent*: the yield should rise in the corridor where substitutes weaken while the stack still works, and the Pressure–Capacity Corridor predicts an inverted U rather than a slope (Section 8.2). The instrument earns its place only if the rate series is published against the macro-state model of

Section 29.2 and the pressure index R of Section 13.7.1, so that convenience yield $= g(R)$ becomes an estimable function rather than a slogan, read as the pre-registered corridor increment of Red Line 16. An increment ruled out is the null result and must be reported as such.

- **Both tails checked.** A convenience yield that appears only in benign states — holders renting out units for yield-farming while nobody needs the bearer service — is a fee-like cash flow, not the regime-contingent service flow, and belongs in the DCF bucket with the rest of Section 13.3.9.

The null result is publishable. The instrument is falsification machinery, not advocacy. If, after the stack delivers measurable service, the lending and basis series show no corridor increment above the pre-registered minimum — holders charge no meaningful premium to part with units in the corridor where substitutes weaken — then the holder-side mechanism is absent at the size the claim needs, and the store-of-value claim has lost its mechanism in public. That outcome belongs on the Value Capture Board beside fee coverage, and it is a more direct test than any red line currently operating on it: the red lines test whether the *chain* broke; this tests whether the *premium* exists.

26.8.7 Buyer Quality and Price Provenance

Wrapper vectors answer *which product* is transmitting exposure. They do not answer *who is absorbing the risk*, with whose capital, under what mandate, at what leverage, against which liability, through which benchmark, and with what capacity to sit through a mark. That is the sovereign-market question of Section 3.1.3, applied to the native asset's own realization plane. Publish, as a checklist rather than as a multiplicative index:

- **Persistence (P).** Does the bid recur because contributions, a mandate, or a liability continue, or because a trade has not yet finished?
- **Horizon (H).** Is the holder matching a multi-decade promise, a quarterly risk limit, or a daily reset?
- **Liability match (M).** Is the purchase offsetting an enduring obligation, or following a weight?
- **Loss-bearing capacity (B).** Can the holder absorb marks without forced sale, margin call, or redemption?
- **Countercyclicality (C).** Does incremental buying increase when prices fall, or does the governing rule absorb less risk precisely then?

A high-persistence, low-match, low-countercyclicality buyer is the automatic index flow of Section 3.1.3: it keeps arriving and still fails as an anchor. A leveraged relative-value fund can score the opposite way on persistence and the same way on loss-bearing capacity. Neither score is a number to multiply into a headline Q_A. The fields exist so that "there is demand"

cannot be reported without saying *what kind*.

Price then requires a provenance model, not only a feed: who bought, with whose capital, under what mandate, at what leverage, with what redemption terms, against which liability, through which benchmark, and with what capacity to survive losses. A policy that treats the resulting price as an external referendum on credibility is reading an instrument the state and the financial industry themselves constructed. That is the epistemic loop Section 3.1.2 already warns about, now stated as a VerifyFlow requirement: **no buyer-quality ledger, no inference from the long yield—and no inference from the native asset's price either.**

26.8.8 Native Monetary Buyer Map

Buyer quality belongs inside the monetary chain, not only in the external price model. Publish a Native Monetary Buyer Map that reconciles protocol flows and holder cohorts without pretending that pseudonymous addresses are known persons:

Just-in-time fee acquisition. Share of fee-paying units acquired within pre-declared windows before expenditure, versus fees paid from retained balances.

Operator sell-through and inventory. Rewards and fee receipts sold, retained, hedged, or pledged by operators, by cohort and interval.

Burns and net issuance. Gross burns, gross issuance, and net issuance separately; a burn is supply arithmetic, not evidence of a buyer.

Self-custodied reserve demand. Net accumulation in addresses or proofs-of-control classified under published, conservative reserve criteria, excluding custodial and wrapper inventories.

Wrapper demand. Custodial, fund, treasury-company, and synthetic exposure reported apart from native reserve holdings.

Leverage. Borrowed funding, derivatives, rehypothecation, liquidation sensitivity, and margin terms where observable.

Holding period. Cohort survival and realized holding-duration distributions, not a single average that mixes fees with reserves.

Loss-bearing capacity. Evidence that the funding source can absorb marks without margin call, redemption, or forced sale.

Countercyclical accumulation. Whether native self-custodied reserve cohorts add units during drawdowns and regime-pressure episodes rather than merely after recoveries, read net of exercise, so that a cohort which spends into a stress episode is not scored as a cohort that left.

Outflow classification. Every reserve-cohort outflow is classed as *forced liquidation* (margin call, redemption, forced sale), *operating turnover*, or *exercise* — units spent on triad services, settlement, or exit during a declared stress episode. Exercise counts for the holder-side mechanism, not against it: a balance that falls because it was used for what it was held for is the mechanism working. Forced liquidation counts against loss-bearing capacity; turnover is neu-

tral.

The map must reconcile identities only to the degree evidence permits. Unknown holdings remain unknown; they are not assigned to the most favorable category. Unknown is a coverage limit, not a quality grade, and taking part in the measurement programme is not a condition of privacy: the map publishes its coverage and selection limits instead. Gross native volume passes Premise 7 only when the retained, self-custodied, unlevered or transparently financed, loss-bearing, countercyclical cohort is material and persistent.

The map also carries a representation discipline that every published aggregate in this chapter inherits: any aggregate over households or holders must disclose its weighting rule. The weighting rule is a moral choice made visible. Dollar-weighted and person-weighted summaries of the same underlying observations can move in opposite directions — and neither is the false one. The arithmetic is exact. The gap between the two means is the covariance of spending weight and experience over mean spending weight, $\bar{x}_{\text{dollar}} - \bar{x}_{\text{person}} = \text{Cov}(C_i, x_i)/\bar{C}$, so a positive covariance between spending weight and experience lifts the dollar-weighted mean above the person-weighted mean; the identity is about means, and a statement about the person *median* needs a separate argument about the shape of the distribution. That is the mechanism behind the "two economies" reading of the same dataset (Section 3.6). Board aggregates therefore publish their weights (units transacted, addresses, identities, delta-equivalent exposure) alongside the number, and where a dollar-weighted and a person-weighted reading of the same series diverge persistently, both are shown rather than the more favorable one. A verified aggregate that describes almost no actual participant is not a measurement failure; it is a choice of weights, and the choice is now visible.

26.8.9 Principal Map

The Native Monetary Buyer Map asks who warehouses native monetary risk. It classifies by behavior and reconciles identities only to the degree evidence permits. It does not ask what kind of thing the holder is, and in an economy of delegating agents (Chapter 7) that omission lets one operator's fleet read as a distributed constituency. The Principal Map is the overlay that closes the gap. It is not a board of its own: it classifies every native flow and every reserve cohort already reported elsewhere by **ultimate principal**, to the degree evidence permits, and publishes the residue as its own category rather than assigning it favorably.

Series published from the map, each on the same window as the board it overlays:
Principal shares of native fee turnover, prover collateral, corridor liquidity, and the self-custodied reserve cohort, by class.
Human-Revocable Share (HRS). The fraction of agent-controlled native balances whose key root is held by a human with standing and whose revocation is a protocol operation. Feeds the ninth requirement (Item 9) and the watch clause on Red Line 12.

Class	Definition	Evidence admitted
Human with standing	A natural person holds the key root, can revoke or cap the acting agent through the protocol, and bears the loss	Proof of control under a human-bound predicate (Section 20.4), with no identity disclosed; revocation path attested
Firm	A legal entity holds the root and bears the loss; agents act under a disclosed policy	Entity credential; policy hash published
Agent policy	The acting process holds the root, or an operator-held root the process can extend; loss falls on an operator balance sheet or nowhere	Policy template hash; operator attestation; absence of a human-revocable root
Unknown	No admissible evidence	None; published as residue, never reassigned

Table 26.5: Principal classes. The first class depends on a Layer 3 predicate that is specified in outline only (Section 20.5); where it is unavailable, holdings that would fall there are reported as *firm* or *unknown*, not assumed human.

Native Policy Concentration Ratio (NPCR). The Policy Concentration Ratio of Section 26.8.3 computed over native positions—collateral, liquidity, reserves—rather than over wrapper flows: the share of each native series governed by the largest policy templates. Exit correlation across templates during declared stress episodes is published beside it, so that if policy concentration proves a poor proxy for behavioral concentration the substitution can be made on evidence.

Agent-originated share by operator class. Fleet-originated versus independently originated native demand, as defined in Section 7.4. Companion to the Homestead Ratio on the Layer 0 Capacity Board.

Cross-Perimeter Share (CPS). The fraction of receipt-bearing hops whose principal chain crosses an operator boundary (Section 7.2). It requires the optional `principal_chain` context field of the PIDL receipt (Appendix A). Until that field is in use across a material share of receipts, CPS is published as *not computable*, with the covered share stated; operator-level inference from settlement graphs is a weaker fallback and is labeled as such wherever it is shown.

Representation discipline. The map inherits the weighting rule of the buyer map: unit-weighted and principal-weighted readings of the same series can diverge, and both are shown. A board that reports ten thousand self-custodied holders and one policy template has reported one holder; the principal-weighted figure says so.

Companion to the triple divergence. The joint failure signature of Section 33.9—positive WNG, no corridor increment in convenience yield, falling Homestead Ratio—gains a companion, not a fourth term: falling HRS. A stack that is losing the closed-stack contest and the wrapper contest while its remaining native activity is agent-controlled with no human root is

failing three ways at once and reading green on the Value Capture Board.

26.9 The Public Boards

Metrics surface as twelve public boards or maps. The original five have grown to twelve as the thesis has extended from pure verification economics into physical capacity, political economy, dependency, holder quality, and market structure:

Proof & Compute Board. Summarizes VERIFYPRICE per workload, SLA attainment, queue depths, and failure modes.

Settlement & Privacy Board. Reports swap success, refund safety by corridor, time-to-finality, and anonymity-set health.

Neutrality & Admission Board. Tracks time-to-first-proof for new entrants, top-N share, geo/ASN distribution, and house share.

Layer 0 Capacity Board. Extends hardware-profile telemetry with Facility Energy Receipts and Facility Capacity Receipts (Section 17.7): grid node, interconnection class, outage history, backup duration, fuel mix, cooling dependency, and Physical VerifyPrice (Section 22.8). Reports both Sovereign Optionality and scenario-specific Delivered Verified Capacity — how much of the grid still lights up when the weather turns — plus active binding constraints, substitution latency, and dependency groups (Section 17.9). Tracks whether the compute/proof layer is physically resilient or just another hyperscaler/cloud bottleneck, including the **Homestead Ratio** (Section 5.6) and, as its demand-side companion, the agent-originated share by operator class and the Cross-Perimeter Share (Section 26.8.9).

Economic Coverage Board. Charts fee-coverage ratio, real-resource VERIFYPRICE against its constitutional SLOs, and demand curves indexed to policy events.

Value Capture Board. Tracks fee coverage, native-asset burn rates, collateral lockup volumes, issuance vs. capacity growth, workload mix (budgeted vs. speculative), and bypass-channel indicators (operator off-ramp rates, stablecoin-denominated fee share), the M1+M2 share of verified units, the link-9 instruments — haircut dispersion, free-choice denomination share with the quote/settle matrix, acceptance breadth — and the off-ledger exposure disclosure series of Section 13.3.10. Native fee turnover and collateral lockup are additionally reported by principal class (Section 26.8.9). This board makes the Value Capture Lemma from Chapter 8 auditable: if value capture is failing, it shows here before it shows in price. It also carries the **triple divergence** composite — persistently positive WNG, no corridor increment in regime-conditioned convenience yield, and falling Homestead Ratio against improving closed-stack service metrics — which is the observable signature of the joint failure mode named in Section 33.9.

Wrapper Dominance Board. Tracks whether the asset is becoming a financial wrapper rather than money: the **Wrapper Dominance Ratio** (Section 13.6) alongside its component series (ET-

F/wrapper AUM, custodied balances, exchange balances, wrapped supply, native fee share, non-custodial settlement volume, shielded settlement volume, collateral lockups). A rising WDR while native usage stagnates is surfaced here before it appears anywhere else.

Agency Preservation Board. Tracks non-custodial usage share, selective-disclosure scope, dossier-reconstitution risk, and forced-disclosure incidence (Section 27.1.2), operationalizing the ninth SoV requirement (Item 9). This board is optional in early phases but becomes load-bearing once the protocol has material institutional and administrative-state interaction. It also carries the Human-Revocable Share of agent-controlled native balances (Section 26.8.9).

Market Realization & Wrapper Board. Reports the VerifyFlow vectors (Section 26.8): wrapper inventory by class, assets and delta-equivalent exposure, creations and redemptions, holder flow elasticity, net mechanical gain, Wrapper Recycling Ratio, constituent concentration, market-depth ratios, hedge-instrument mix, counterparty concentration, the Wrapper–Native Growth Gap, custodial versus self-custody shares, recursive wrapper exposure, cross-jurisdiction exposure, holder-cohort statistics, and the convenience-yield telemetry of Section 26.8.6 (lending rates, forward and perpetual basis, wrapper basis, published against the regime-pressure index).

Native Monetary Buyer Map. Separates just-in-time fee acquisition, operator sell-through and inventory, burns and net issuance, self-custodied reserve demand, wrapper demand, leverage, holding period, loss-bearing capacity, and countercyclical accumulation (Section 26.8.8). It asks who warehouses native monetary risk, not merely whether demand exists. Every cohort on it is read through the Principal Map overlay (Section 26.8.9)—principal shares and the Native Policy Concentration Ratio—so that one operator's fleet cannot read as a constituency.

Common-Cause Dependency Board. Maps dependencies crossing nominally independent facilities or jurisdictions: hardware and accelerator vendors, firmware and signing systems, energy and fuel markets, cables and network routes, identity, custody, market making, banking, legal recognition, cooling-water basins, and software distribution. For each dependency publish affected capacity, substitutable alternatives, substitution latency, correlated stress scenarios, and the DVC loss if the group fails. The correlated physical-failure framing is adapted selectively from Doomberg's *A Minor Detail* [Doomberg 2026c], treated as scenario commentary rather than forecast.

Sponsor Dependency Board. Reports subsidy concentration, politically contingent demand, cloud-credit dependence, dominant-customer exposure, state or foundation veto power, operator exit capacity, sponsor-withdrawal sensitivity, conflicting terminal objectives, capacity removable on short notice, and spillover to neutral third parties. Subsidized operators and sponsors are separate strategic actors even when their current incentives align. The principal–agent framing is adapted selectively from Doomberg's *Over Target* [Doomberg 2026d]; inferred intent and tactical escalation claims are excluded.

Do not merge the Value Capture Board and the Market Realization & Wrapper Board.

They answer different questions, and collapsing them destroys the distinction the thesis depends on:

- **Value Capture Board:** *Is service value accruing to the native asset?*
- **Market Realization & Wrapper Board:** *What is currently setting its price?*

The Value Capture Board stays focused on protocol economics. The Native Monetary Buyer Map asks whether a monetary-risk warehouse exists. VerifyFlow stays focused on market realization. A reader needs all three to distinguish service accrual, monetary adoption, and financialization.

These boards are not marketing pages; they are part of the protocol's public interface.

26.10 Detecting Drift

Drift is the failure mode that never announces itself. Nothing breaks; everything erodes — a few milliseconds added to p95 here, a few points of share consolidated there — and each reading is individually explicable, which is precisely why an adversary (or mere entropy) prefers it to an attack. Telemetry exists to catch the trend while it is still a trend:

Verification cost creep. VERIFYPRICE shows p95 verify time creeping up. Response: R&D lab proposes fixes; governance ratifies changes.

Verifier concentration. 80–90% of verifications performed by one hardware profile. Response: adjust rewards to subsidize diverse profiles.

Reachability degradation. VerifyReach shows a major country blocking protocol ports. Response: shift transports, promote alt routes. You can dam a river; you cannot dam the rain — the response is always to add paths faster than the adversary can close them.

Corridor fragility. VERIFYSETTLE shows corridor success dropping. Response: re-weight corridors, onboard new LPs.

26.11 Making Neutrality and Repression-Resilience Falsifiable

Neutrality and repression-resilience are usually sold as moral properties. In this thesis, they are **falsifiable hypotheses**:

Neutrality hypothesis: "No actor or group can reliably censor, privilege, or front-run specific flows across Layers 1–5 without being detected and without other actors being able to route around them."

Repression-resilience hypothesis: "Under YCC, capital controls, DPI, and blacklists, the triad's core capacities retain sufficient DVC and the base asset's self-custodied holder constituency retains sufficient agency to remain a monetary candidate."

Telemetry makes these hypotheses testable:

- If VerifyReach shows certain flows consistently blocked and no alternatives emerge, neu-

trality is falsified.
- If VerifySettle shows privacy corridors collapsing under regulation, repression-resilience is falsified.
- If VerifyPrice shows verification cost pushed beyond reach, "anyone can verify" is falsified.

A system that cannot tolerate bad news cannot be money; it can only be marketing.

26.12 Policy-Attack Stress Harness

A monetary stack designed for repression should publish the exact conditions under which it considers itself passing or failing. For each attack scenario, the tests and pass criteria are:

Yield-curve control & negative real yields:
- **Test:** Over 24–36 months of -300 to -500 bps real yields, FeeCoverage (retained fees, burns excluded) as a share of security budget is stable or rising.
- **Pass:** Tokens behave like claims on useful capacity, not synthetic bonds.

On-/off-ramp throttling:
- **Test:** Non-custodial corridors maintain $\geq 95\%$ success with no protocol-attributable loss (Section 23.4); per-asset anonymity metrics remain above their thresholds (Section 23.5).
- **Pass:** Users can move capital via non-custodial rails even when custodial ramps are hostile.

Provenance mandates:
- **Test:** Canonical provenance workloads have healthy VerifyPrice; regulators can run `verify(receipt)` with commodity tooling.
- **Pass:** The triad is the cheapest way to comply with provenance mandates.

Hardware backdoor mandates:
- **Test:** A non-trivial share of receipts are tied to open or sampled hardware profiles, not a single opaque TEE.
- **Pass:** Hardware mandates become one option among many, not a kill switch.

When the stack claims "we can survive policy attacks," the proof is not a blog post; it is a corpus of **stress-harness runs with metrics and receipts**.

26.12.1 Market-Realization Stress Tests

The harness above stresses repression and availability. The Market Realization Plane needs its own scenarios, because the compositional adversary (Section 5.2) does not respond to any of the tests above. Each scenario is run against the VerifyFlow vectors of Section 26.8.

A. Wrapper inflow cannon. A large spot or thematic fund launches; creations arrive independent of daily returns; holdings are concentrated; native liquidity is shallow. **Measure:** Q^{RD}/V, constituent concentration, price impact, native-use response, WNG. **Question:** does the bid convert into native usage, or only into price?

B. Elasticity regime break. A $2\times$ or $3\times$ wrapper grows rapidly and holder elasticity drifts toward zero, so recycling declines. **Measure:** κ, Wrapper Recycling Ratio, volatility feedback, reversal symmetry. **Question:** has a product quietly become a momentum amplifier?

C. Creation halt. Wrapper inflows stop while protocol-native usage stays constant, removing passive support. **Measure:** price drawdown, native fee coverage, governance response. **Pass:** native fee coverage remains sound *and* governance does not respond by inflating issuance or subsidizing price.

D. Overnight gap or market halt. The underlying gaps beyond a daily-reset wrapper's viable range before it can rebalance. **Measure:** wrapper solvency, dealer collateral calls, liquidation paths, cross-venue settlement. **Pass:** native protocol operations continue unaffected; wrapper failure stays contained in the Market Realization Plane.

E. Swap-capacity exhaustion. Dealers reduce swap exposure and funds migrate to options. **Measure:** implied-volatility demand, gamma hedging, dealer concentration, collateral requirements, spot–volatility correlation.

F. Price up, protocol flat. Price triples and wrapper AUM grows while native fees, settlement, proofs, and compute usage do not. **Expected result:** market success, monetary-adoption failure, and an *automatic* dashboard warning. A thesis that cannot raise an alarm during a rally is not falsifiable.

Chapter 27

Legal, Policy, and Jurisdictional Posture

The adversary here does not send exploits; it sends letters. Layer 6 sits where protocol meets law and politics, and the encounter is asymmetric: code can be amended in a release cycle, but sanctions regimes, travel rules, and court precedents move on their own clock and do not consult the changelog. If the triad is to be sustainable, it must support **lawful privacy**, defend itself against bans without being defined by them, and avoid capture by any single jurisdiction. That means designing for the law as it is — including where it is hostile — rather than for the law as the cryptography would prefer it.

27.1 Lawful Privacy as Protocol Design

"Lawful privacy" is often a euphemism. Here it means something precise:

- **By default**, settlement and compute flows are private.
- **By design**, participants can opt into **selective disclosure** via viewing keys and receipts.

Constitutional Constraint: Policy = Predicates

Policy compliance is defined as **ZK-checkable predicates** over credentials and receipts, plus **selectively disclosed audit windows** via viewing keys.
It is **not** defined as global graph inspection or universal traceability. Any policy requiring ubiquitous transaction tracing is treated as **incompatible with the monetary design**.

27.1.1 What Lawful Privacy Is Not

Lawful privacy is not a promise that users can ignore legal obligations. It is a design posture in which the protocol minimizes unnecessary disclosure while allowing users to produce scoped, auditable evidence when they choose or are legally required to do so. Specifically:

- Lawful privacy does **not** guarantee immunity from subpoena, tax reporting, sanctions enforcement, or employment obligations.

- It does **not** mean that regulators will accept predicate-only compliance in all jurisdictions. Some jurisdictions will reject it.
- It does **not** mean that users can evade capital controls or hide from governments. The protocol provides *technical* confidentiality, not *legal* immunity.
- It **does** mean that the protocol itself does not require surveillance infrastructure, and that disclosure is always at the holder's discretion rather than built into the protocol.

The legal posture is **defensive, not antagonistic**: the thesis designs for resilience under legal heterogeneity, not for confrontation with law.

27.1.2 Social Coercion Metrics: When "Optional" Becomes Mandatory

The protocol already acknowledges that lawful privacy is a technical guarantee, not a social guarantee (Section 5.6): it cannot prevent coercion to disclose keys, but it can raise the cost of mass surveillance, preserve optionality, and make coercion visible through telemetry. The remaining question is *when* optional disclosure becomes mandatory in practice.

Forced Disclosure Incidence. Reported cases where employers, banks, landlords, visa processes, exchanges, or benefits agencies require viewing keys.

Selective Disclosure Scope. Average breadth of information disclosed per audit event.

Dossier Reconstitution Risk. Probability that repeated selective disclosures can reconstruct the user's full graph.

Non-Custodial Practicality Index. Share of flows that can be completed without a KYC chokepoint at equivalent service quality.

Lawful privacy fails socially before it fails cryptographically. If viewing keys become routine prerequisites for employment, banking, housing, travel, or benefits, optional disclosure becomes administrative disclosure—the social-layer expression of the Administrative Repression pattern introduced in Section 5.2. The protocol cannot stop all coercion, but it can measure whether coercion is becoming systemic. These four metrics feed the Agency Preservation Board (Section 26.9).

Read the table honestly. It is tempting to head the right column of this table "Cannot Require," as if the protocol held a boundary that regulators must respect. That framing is wrong. The Travel Rule (FinCEN; FATF Recommendation 16 [FATF 2025]; the EU Transfer of Funds Regulation) obliges obligated entities to *transmit and retain actual originator and beneficiary identity data* — name, account, amount — not set-membership proofs. A regulated institution cannot lawfully satisfy that obligation with predicates, no matter how sound the cryptography. The accurate statements are:

- **Within the current law, predicate compliance is an engineering option that is not yet a legal option.** It requires legislative or rule-level change in every major jurisdiction before a regulated counterparty can rely on it. The protocol's contribution is to make that change

Requirement	How Satisfied (Protocol)	What Current Law Actually Requires
AML screening	ZK proof of "sender $\in$ cleared-set"	Originator identity data (Travel Rule)
Jurisdictional limits	ZK range proof of amount	Exact amount above reporting thresholds
Tax audit	Epoch-scoped viewing keys over rotated addresses (Section 23.7)	Records sufficient to reconstruct transactions
Counterparty verification	ZK set-membership proof	Named-party identification

Table 27.1: Lawful privacy: what the protocol can satisfy versus what current law demands. The right-hand column is not a design boundary the protocol can hold by engineering; it is the law as written, and predicate compliance is *unlawful* under it in most major jurisdictions until the law changes.

expressible: when a jurisdiction moves, the machinery exists to comply with far less surveillance than universal tracing would need.

- **Outside regulated perimeters, the predicates are simply private.** P2P settlement between non-obligated parties is not a Travel Rule event anywhere today. The design buys lawful headroom at the edges while the institutional perimeter remains a policy question, not a protocol one.
- **Who curates the cleared set is a hard, unresolved question.** A "sender $\in$ cleared-set" proof requires an oracle maintaining the set: a sanctions-list curator inside a privacy protocol is simultaneously a chokepoint, a legal attack surface, and a party with liability. The thesis does not have an answer and declines to pretend otherwise.

27.2 Securities Law: Classification Is Transaction-Specific

The strongest legal threat to the asset is not a ban. It is classification, and part of the exposure is created by the thesis's own design choices, so Chapter 28 cannot treat classification as weather. This section is dated: its sources are as of 13 October 2026, and the posture moves.

The governing frame. The test remains *Howey* [SEC v. Howey 1946] — an investment of money, in a common enterprise, with a reasonable expectation of profit derived from the efforts of others — and the Commission's interpretive release of March 2026 [SEC 2026a] states expressly that it does not supersede or replace that test (§I). The release separates two questions that are easy to run together: whether a crypto asset is itself a security, and whether a transaction in it is subject to an investment contract. A non-security asset can be sold subject to an investment contract; the promises and circumstances of the sale decide whether it is; and an issuer that fails to perform the managerial efforts it promised can face liability for that offering even after the asset has separated from it (§IV). Governance rights do not by themselves make

an asset a security: the release contemplates digital commodities whose holders vote on technical and governance matters such as software upgrades and treasury expenditures (§III.A). Protocol mining and protocol staking are addressed by the Commission itself (§V.A, §V.B), so this analysis rests on the release rather than on the earlier Division staff statements on mining and staking [SEC 2025a] [SEC 2025b], which have no legal force.

A transaction inventory, not a token label. The design creates several kinds of transaction, and each is examined on its own facts: what is promised, who controls performance, what financial rights exist, and what the purchaser is asked to rely on.

Transaction	What is promised, and by whom	What the purchaser relies on	The open question
Native issuance to provers	Rewards on a schedule fixed ex ante, by protocol rules	The schedule and the provers' own work	How useful-work consensus compares with the protocol-mining facts of §V.A, given that workload selection and client demand enter where a hash puzzle has none
Secondary transfers of the base asset	Nothing, by anyone	Market liquidity	Whether any party's continuing representations — a burn programme, a roadmap — attach to the asset after launch
Service prepayments (Work Credits)	Reserved service, by a named provider (Section 12.1.1)	That provider's delivery, with refund on stated terms	A consumptive contract on its face; closer to an investment if marketed for resale value or if proceeds finance construction
Staking and operator pools	Rewards for operating work, with slashing	The operator's work, or a pool operator's	Where pooled arrangements with a central operator fall relative to the protocol-staking facts of §V.B
Project finance	Repayment or return on capacity not yet built	The developer's efforts	Securities-like by design; labelled, offered, and disclosed as such (Section 34.5.2)

Table 27.2: The design's transaction inventory for classification analysis, as of 13 October 2026. Questions to examine, not conclusions.

Burns and fee routing. The fee split's burn is where the economics and the law meet most directly. Economically, the thesis treats a burn as a buyback (Section 13.3.9). On the legal side, the Division of Corporation Finance's staff FAQ, updated 28 September 2026 [SEC Division of Corporation Finance 2026b], addresses buyback programmes, protocol-funded burns included (Q2.5): where a crypto system is functional and has no central party, announcing a buyback of a non-security asset is not a representation or promise to undertake essential managerial efforts, while for a system that is not functional, presenting a buyback as creating yield or return could be. The FAQ is a staff view without legal force, and its qualification — *no central party* — is the one this design has to earn. A burn fraction or fee split that token governance can adjust is discretion over the return to holders. Whether governance concentrated enough to exercise it amounts to a central party is the question to examine, not a settled answer in either direction. The same caution runs the other way: a fixed supply schedule is not legal clearance. The rights, the control structure, the representations, and the transaction decide together.

The design tension, as a question. Removing holder-returning mechanics (no burns, purely consumptive fees) weakens Condition 2 of the lemma and moves the asset toward a pure service medium. Keeping them strengthens the cash-flow claim and keeps governance concentration and discretionary return management open as legal questions. Fixing issuance and fee policy *ex ante*, beyond governance reach, removes the discretion that would make anyone's continuing efforts the source of return, and it is the version closest to what a store-of-value constitution should look like anyway (Section 25.6); it narrows the exposure without conferring clearance. Red Line 7 tracks whether lawful participation remains possible; the classification of each transaction in Table 27.2 determines whether it can be satisfied in the largest capital pools.

This section is analysis, not advice. Jurisdictions differ, enforcement posture shifts, and no statement here is legal advice. Any part of this architecture needs design-specific legal review before it is offered to anyone. The thesis exposes the legal questions accurately; it does not decide them through token terminology.

27.3 Defense-in-Depth Against Bans and Sanctions

Regimes may attempt to ban privacy assets, sanction addresses, mandate KYC at all edges, or classify triad services as "unlicensed financial institutions."

The live precedent for this section is not a statute but a sanctions action and the prosecutions that followed it, and the sequence has to be read to its end, because it is tempting to stop at the point most favourable to the thesis. OFAC designated Tornado Cash in August

2022 [OFAC 2022], treating a non-custodial, autonomously operating privacy protocol *as such* as a sanctionable entity. In *Van Loon v. Department of the Treasury* [Fifth Circuit 2024] (5th Cir., November 2024) the Fifth Circuit held that the protocol's immutable smart contracts are not "property" of a foreign national and therefore not designable under IEEPA, and OFAC delisted the contracts on 21 March 2025 [OFAC 2025]. That is the half usually told. The other half: Alexey Pertsev was convicted of money laundering in the Netherlands in May 2024 [Rechtbank Oost-Brabant 2024]; and Roman Storm was convicted in the Southern District of New York on 6 August 2025 of conspiracy to operate an unlicensed money-transmitting business (18 U.S.C. §1960), with the jury hung on the money-laundering and sanctions-conspiracy counts (*United States v. Storm* [SDNY 2025]). The lesson is therefore not the one the first half suggests. **Non-custodial design defeated designation of the *code*. It did not defeat prosecution of the *coders*.** A court agreed that nobody controlled the contracts, and a jury nonetheless convicted the person who wrote and maintained the front end of running a money-transmitting business. Developer liability is not a residual risk after the architecture is right; it is a first-class risk surface that the architecture does not reach. Three consequences for this stack:

- **Developer and operator liability is a first-class risk surface**, not a footnote to Layer 5's corridor design, and it is the surface that survived the strongest possible architectural defence. Anonymous or distributed development, jurisdictional separation of maintainers from operators, the absence of any operated front end or relayer that a prosecutor can call a "business," and legal entity structure (Section 27.4) are not conveniences; they are the difference between a delistable product and a prosecution target, and after *Storm* they are not a guarantee of either.
- **Exchange delisting is the working enforcement mechanism, not the ban itself.** Monero and Zcash were not outlawed in most jurisdictions; in the jurisdictions where regulated venues operate under Travel Rule and MiCA-style obligations [EU 2023] — the EU, the UK, Japan, South Korea, Australia, and the large US venues — they were made practically inaccessible by compliance burden at those venues, while remaining freely tradable elsewhere and peer-to-peer everywhere. Corridor design must assume the fiat on-ramp is the chokepoint that closes first (Section 13.5.1 engages the substitute channel this creates).
- **Non-custodial is a legal argument against designation, not a shield against prosecution.** The strongest available defence of the *protocol* is precisely the non-custodial architecture this thesis specifies — no controller, no admin keys, no ability to freeze — because it is the fact pattern that defeated the Tornado Cash designation in court. The design should maximize those properties deliberately, because they are load-bearing in litigation over the code. They did nothing for the coders, and the thesis does not pretend they will do more for this stack's.

Layer 6 cannot prevent law from existing, but it can:

- Avoid giving any single jurisdiction a **kill switch** — including the one the protocol's own emergency path would otherwise hand to whoever can compel three of five signers, which

is why the censoring subset of that path is advisory by default (Section 25.5.3).

- Maintain **jurisdictional diversity** in hardware profiles, proof factories, corridor LPs, and foundation incorporation.
- Provide **fallback modes**: local-first clients, mesh/satellite relays, minimal CLI modes that maintain basic functionality under degraded conditions.

The legal posture is **resilience under legal heterogeneity**: the stack does not advertise itself as a tool to evade law, but designs for continued operation across a diverse legal landscape where some jurisdictions are friendly, some hostile, and most ambiguous.

27.4 Labs, Foundations, and Neutral Router Commitments

Key commitments:

Neutral router charters. Router operators commit to content-agnostic routing. Deviation is punished in-protocol and reputationally.

Foundation / lab structure. Legally separate entities with public charters that define roles and explicitly renounce certain powers.

Multi-jurisdiction footprint. Incorporations and key staff spread across legal regimes to avoid single-point capture.

Chapter 28

Operator & Investor Checklist

The last piece of Layer 6 is a practical checklist: how builders, operators, and allocators decide whether this stack qualifies as a next-generation store of value. It is deliberately unglamorous, and that is the point. Everything before this chapter argued the thesis at full strength; a checklist is where you stop believing the argument and start checking it. Read it as the stack's own answer to "what would change your mind?" — each question below has a yes that strengthens the claim and a no that weakens or falsifies it.

28.1 Triad Supply: Privacy, Proofs, Compute

For each leg of the triad:

Privacy:

- Are there live, non-custodial privacy corridors with measurable anonymity sets?
- Are shielded pools growing in depth and churn?
- Are lawful-privacy patterns (viewing keys + receipts) actually used?

Proofs:

- Is there robust supply of proofs for canonical workloads?
- Are VERIFYPRICE SLOs met consistently?
- Is proof production diversified across hardware profiles and geographies?

Compute:

- Are meaningful amounts of useful compute being proven?
- Is there a liquid market for Work Credits?
- Are triad workloads tied to real-world demand?

28.2 Value Capture & Anti-Bypass

The most important questions for evaluating whether the SoV thesis holds:

- Can users buy equivalent Privacy, Proofs, or Compute service **without the native asset**?

- Are fees **actually paid in the asset**, or do most operators accept and immediately off-ramp alternatives?
- Are burns and collateral lockups **material** relative to total supply?
- Is there a visible **bypass channel** (hyperscaler-hosted provers, stablecoin-denominated services) that captures most of the demand without flowing through the asset?
- Does governance have the power to **redirect value away from holders** (e.g., by eliminating burns, inflating supply, or changing fee routing)?
- Is the **Wrapper Dominance Ratio** (Section 13.6) rising while native fee share, private settlement, and collateral lockups stagnate?
- Is **price** being cited as evidence of monetary adoption in place of native fee, burn, collateral, and settlement series (Corollary 13.1)?

If the answer to any of these is "yes," the system may be useful infrastructure but the SoV thesis is weakened or falsified.

28.3 Institutional Holdability

- Is **custody support** available from regulated custodians?
- Is the **legal classification** clear enough for treasury or fund mandates?
- Is the asset **auditable** enough for institutional reporting requirements?
- Is **liquidity** deep enough for institutional-scale positions?
- Is there a **valuation policy** that allocators can apply (fee-flow models, comparable frameworks)?

Holdability is a virtue with a bill attached (Section 5.6). Every question above should therefore be paired with its market-realization counterweight:

- Does wrapper growth **seed** native usage, or **substitute** for it? Is the Wrapper–Native Growth Gap (Section 26.8) persistently positive?
- What share of supply sits under custodial control (CCR), and how concentrated is that custody?
- Is delta-adjusted synthetic exposure (SER) large relative to liquid free float?
- Do daily-reset leveraged products exist on the asset, and what is their net mechanical gain κ and Wrapper Recycling Ratio?
- Is exposure concentrated in a small number of dealers or counterparties, and has the hedge-instrument mix shifted from physical to swaps to options?
- Is demand governed by a small number of common rebalancing policies (Policy Concentration Ratio)?
- Where native demand or native holdings are agent-operated, are they classified by principal (Section 26.8.9)? What share is one operator's fleet, what share is governed by the largest policy templates on the *native* side, and what share of agent-controlled balances can a person

still revoke through the protocol?

28.4 Holders and the Monetary Mechanisms

Everything above asks whether the conditions for moneyness hold. These questions ask whether it is happening, and fees, burns, and collateral cannot answer them.

- Does a persistent, self-custodied, loss-bearing cohort add through drawdowns (net of units spent on the service during stress, which count for it), or is native demand acquired just in time, sold on receipt, leveraged, or wrapped (Section 26.8.8)?
- **Why hold rather than convert on the day?** Has the access experiment shown a material advantage for a held balance over the best Bitcoin-funded strategy, including provider-side conversion, for any workload and stress state (Section 13.1)?
- Is the persistent bid a rule that absorbs less risk as prices fall, or leverage that leaves on a margin call (Section 3.1.3)?
- Where agents hold the asset, does anyone ever revoke them? A human-revocable share that holds while the policy concentration of the same cohort rises is standing exercised by nobody (Section 5.2).
- **Would a lender accept it without investigating it?** Is cross-lender haircut dispersion on verifiable positions, across independently maintained underwriting approaches, lower than on opaque comparables of the same volatility (Red Line 17)?
- **Does anyone reckon in it?** Is free-choice denomination growing, and do insiders denominate their own short-tenor obligations in the unit (Red Line 18)?
- **Is holding it worth more when substitutes fail?** Does the convenience yield rise in the corridor band above the pre-registered minimum increment (Red Line 16)?

28.5 Stack Health: Layers 0–6

Layer 0. Documented hardware profiles? Open hardware participation? Visible incident reports? Facility Capacity Receipts available and independently verifiable (Section 17.7)? Physical VerifyPrice within bounds (Section 22.8)? Issuance bounded by stressed Delivered Verified Capacity, with the active binding constraints named (Section 17.9)?

Layer 1. VerifyReach shows resilient connectivity? Alternative transports exist?

Layer 2. Clients distributed via multiple channels? Reproducible and signed binaries?

Layer 3. Pseudonymous credentials integrated? Actors can prove rights without doxxing? Non-custodial usage share and dossier-minimization score visible (Item 9)?

Layer 4. Canonical workloads documented? VERIFYPRICE within SLOs? Proof production decentralized?

Layer 5. Non-custodial corridors healthy and diversified? Refund safety enforced?

Layer 6. Public dashboards? Documented governance processes? Recent incident reports?
Market Realization Plane. Is VerifyFlow published at all (Section 26.8)? Are wrapper inventory, creations/redemptions, flow elasticity, concentration, and the Wrapper–Native Growth Gap visible? Is the market-realization stress harness (Section 26.12.1) run and reported? Remember that this plane sits *around* the stack, not inside it: its failure does not break the protocol, but its opacity breaks price interpretation.

28.6 Telemetry Honesty

Telemetry itself can be gamed. Checklist for honesty:

- **Open data.** Are raw metrics public? Can independent teams reproduce Verify* metrics?
- **Client diversity.** Metrics gathered by multiple operators?
- **Incentive alignment.** Any incentive to under-report problems?
- **History.** Historical series for metrics, or only recent snapshots?

28.7 Red Flags and Failure Patterns

Finally, some patterns that should trigger skepticism:
Closed hardware monoculture. One TEE, one vendor, no sampling, no open profiles → Layer-0 cliff.
Opaque bridges. "Magic multisig" bridges with no proofs, unclear jurisdiction, and no incident history → not settlement rails, just custodial risk.
Foundation fiat. Major parameter or policy changes via blog post, with no on-chain trace, no telemetry, and no incident report → governance capture.
Zombie corridors. Privacy rails that have not moved meaningful volume in months but still appear in marketing. Anonymity sets that are effectively dead.
Proof theater. Lots of "ZK" branding but no public VERIFYPRICE metrics, no canonical workloads, no commodity verifiers.
Governance theater. Token votes with single-digit participation deciding fundamental parameters; no constraints; no SLOs.
Price as proof. Marketing that cites market capitalization, wrapper AUM, product launches, or trailing returns as evidence that the monetary thesis is working, while native fee, burn, collateral, and settlement series are absent or flat → financialization narrated as adoption (Corollary 13.1).

Any one of these is survivable in the short run; taken together, they say:

> *"This is not a triad-backed money system; it is a platform dressed as one."*

A system with **no telemetry** is untrustworthy. A system with **one vendor's telemetry** is fragile. A system with **multi-source, reproducible telemetry** has a shot at being money.

Part V closes the governance loop:

- It treats SLOs and dashboards as the **constitution**.
- It makes **Verify* telemetry** the lifeblood of neutrality and repression-resilience.
- It frames law and politics as **constraints to be designed for**, not as deities to be appeased.
- It gives builders and allocators a **practical checklist** for deciding whether a system's "store of value" claim is mathematically grounded or just well-typeset.

With Layer 6 in place, the stack has all three angles: Money (triad as SoV), Stack (Layers 0–6), and Telemetry (Verify* keeping it honest). Part VI can now focus on dynamics (adoption curves, risk, and implementation sketches): how this thing actually gets built, traded, attacked, and, if it works, quietly becomes part of what the world calls "money."

Part VI
Dynamics, Risk & Implementation

Parts I–V defined the triad as distinct services that may support a conditional base-asset monetary candidate, named the delegating machine increasingly buying them, built a seven-layer stack to supply them, and wired in governance and telemetry so deliverability, holder quality, neutrality, and repression-resilience are falsifiable. Part VI asks the unromantic question:

> *What happens when this stack hits the real world? Who uses it first, how does it fail, and under what conditions does the separate base asset remain a credible monetary candidate?*

Chapter 29

Adoption Curve & Ecosystem Dynamics

The stack is not adopted by "the market" in one step. It is pulled in by different constituencies, each with their own pain.

29.1 The Transition Window: Hard Assets as Bridge

Before the triad stack is operational, rational actors face a dilemma: soft guarantees are visibly failing, but the new infrastructure is not yet robust enough to absorb serious capital. Gold and silver—time-tested, non-digital, and harder to seize in physical form than any account balance, though not immune to it—remain sensible bridge assets during this window. Executive Order 6102 in 1933 and Australia's 1959 gold-delivery requirement were compulsory-delivery requirements on physical gold, imposed by statute on holders who had done nothing but hold it.

This thesis does not argue against holding them. It argues that once VerifyPrice, VerifyReach, VerifySettle, and DVC are live and healthy, the triad stack may offer distinctive digital services:

- **Verifiability:** Gold's purity requires assay; proofs verify in milliseconds on a laptop.
- **Programmability:** Gold cannot enforce conditional logic; typed Work Credits can encode service SLAs, escrow, and policy without becoming money.
- **Portability under surveillance:** Physical gold is detectable at borders; shielded settlements are not.
- **Divisibility and settlement speed:** Gold settles in days with intermediaries; privacy corridors settle within an hour or two without custody (deployed BTC↔XMR swaps take 30–90 minutes end to end, bounded by confirmation depth on both chains).
- **Service accrual:** Triad usage can route fees and burns to the associated base asset, producing a cash-flow claim and floor rather than proving monetary premium.

The question is not *whether* to hold hard assets, but *when* the new stack becomes robust enough to migrate. The adoption curve below provides the metrics: once Phase I gate condi-

tions are met (VerifyPrice dashboards live, corridors refund-safe, receipt volume non-trivial), the rational rebalancing can begin. Until then, gold and silver remain what they have always been—stores of value that do not require trust in any single institution, even if they lack the programmability and verification properties of the triad.

> *Gold is the bridge; the triad is a candidate destination if telemetry and adoption gates are met. The bridge is still load-bearing while the destination is under construction.*

29.1.1 Portfolio Migration: From Bridge Assets to Triad Assets

Migration should be telemetry-gated, not narrative-gated. Gold, silver, Bitcoin, front-end liquidity, and physical infrastructure remain the bridge while the destination is under construction. Triad exposure should scale only as the stack clears observable gates:

Gate 0 — Bridge regime. Hold old-world hard assets and liquidity; triad assets are venture/research exposure.

Gate 1 — Verification. VERIFYPRICE dashboards are live for canonical workloads; p95 verification cost remains within SLOs across multiple quarters.

Gate 2 — Settlement. VERIFYSETTLE shows non-custodial corridors maintaining high success and refund safety under stress.

Gate 3 — Reachability. VERIFYREACH shows access across jurisdictions, ASNs, and censorship regimes.

Gate 4 — Value capture. Native-asset fee share, burns, collateral lockups, and non-bypassability are material and publicly visible.

Gate 5 — Institutional holdability. Legal classification, custody, accounting, liquidity, and operational controls allow non-speculative treasury allocation.

Gate 6 — Repression stress. Under a real or simulated repression episode, the stack remains usable and FeeCoverage (retained fees, burns excluded) is stable or rising.

Before Gate 1, the asset is research exposure. After Gate 3, it is infrastructure exposure. After Gate 4 and Gate 5, it becomes a credible store-of-value candidate. After Gate 6, it begins to justify reserve-asset language.

This does not mean hard assets become obsolete once the triad matures. Gold requires no network, no software, no telemetry, no cryptographic agility, and no functioning power grid; it remains rational non-digital redundancy even in a mature-triad world, not merely a placeholder for it.

How the window closes. The bridge-then-destination picture can be read as a waiting problem: hold gold, watch the gates, migrate when telemetry clears. The capital-cycle framing above (Section 29.2) suggests the window is more likely to close by event than by accrual. Historically, the installation phase of an infrastructure buildout ends not when adoption grad-

ually saturates but when its financing structures — obligations written against future utilization assumptions — are repriced in a clearing event. The installed base then changes hands from financially weak owners to stronger users at lower capital cost. The infrastructure does not vanish; it gets repriced [Perez 2002]. For this thesis the implication is a specific one and cuts both ways. The demand base the triad's fee instruments measure — verified compute purchased as a line item by ordinary budgets rather than by frontier labs — is disproportionately a *post*-repricing phenomenon: cheap installed capacity is what makes verification a routine procurement decision rather than a strategic bet. A clearing event that repriced AI-era compute capacity would arrive *before* the fee base it would eventually feed. The same event would also stress every other instrument simultaneously, which is why the phase gates below are metric-anchored rather than calendar-anchored, and why Gate 6 tests the stack through a stress episode rather than during calm. The thesis does not predict the event. It notes that the bridge metaphor's unstated assumption — an orderly handoff — is the historically unusual case, and that the gates are designed to be readable through a disorderly one.

29.2 Macro Policy State Model

AfterFiat is a structural architecture, not a tactical Treasury short. The same sovereign system can pass through different states, and policy repair can extend incumbent arrangements for years. The model therefore specifies states and observables, not a dated sequence:

The market-repair distinction is adapted selectively from Michael Green's commentary in *What the Treasury Needs* [Green 2026c], and the observable repair sequence and its falsifiers from *Sometimes, You Get What You Need* [Green 2026d]. The duration-buyer tripwire is the automatic-buyer argument of Section 3.1.3. They supply mechanisms and scenarios, not adopted quantities or tactical forecasts.

State 1 — Market repair succeeds. Long yields are driven principally by term premium and impaired duration absorption while long-dated inflation compensation remains contained. Buybacks, maturity changes, regulatory repair, or lower front-end rates restore market function; real rates fall, mortgage duration can shorten, and private substitutes remain credible. The immediate repression premium for an open bearer stack may narrow even while the secular agency case remains.

State 2 — Physical inflation blocks repair. Energy, refining, gas, power, shipping, cooling, transformer, or industrial bottlenecks lift long-dated inflation compensation. The financial system may need easier policy while the physical system argues for restraint. Layer 0 costs rise at the same time demand for neutral assets may rise. Long-dated breakevens and inflation swaps are the tripwire distinguishing this state from a pure term-premium repair.

State 3 — Market repair fails and repression follows. Liquidity support, issuance changes, or easing do not restore a durable marginal buyer, and the state shifts from market repair toward

captive demand, preferential regulation, negative real returns, capital controls, or restricted exit. This is financial repression, not merely intervention.

State 4 — Competent closed-stack stabilization. The state successfully coordinates energy, industry, compute, identity, payments, and duration warehousing. Functionality and collateral stability improve while user custody and practical exit contract. This state is not collapse; it is the strongest strategic competitor to the open stack.

Transitions are conditional. Quiet long-dated inflation compensation keeps State 1 available; a sustained rise associated with physical bottlenecks moves the system toward State 2. Failed repair plus coercive balance-sheet rules indicates State 3. Demonstrated closed-stack service delivery with shrinking user exit indicates State 4. None of these observations supplies a tactical forecast for bonds, currencies, or commodities.

A cross-cutting observable: household transmission. The four states are sovereign-level descriptions, and each has a household-level counterpart that can diverge from it — a state can be stable at the sovereign level while its households deplete, and the divergence is a transition signal in its own right (Section 3.6). The observables, all public series: the saving rate; delinquency by income cohort rather than in aggregate; first-time-buyer share and payment-to-income for *new* entrants against incumbents; underemployment and degree-required employment share; mean versus median net worth; and sentiment against spending (a persistent gap between dollar-weighted activity and person-weighted assessment is the divergence itself, not noise). The reason this belongs in the state model rather than only in the adversary analysis of Chapter 5: the states describe who receives sovereignty, and household depletion is how a state loses the political capacity to remain in its current state. A State 1 with exhausted household buffers transitions differently than a State 1 with intact ones — toward whichever policy regime the depleted median will vote for, which is rarely the market-repair state.

A cross-cutting clock: the capital cycle. The four states describe the sovereign system's condition; they say nothing about *when* within the buildout of the new infrastructure that condition is being read. Perez's technological-surge sequence — installation (irruption, then Frenzy), a Turning Point, then deployment (synergy, then maturity) — supplies the missing axis, and it is a clock rather than a state because it advances independently of which state the sovereign system occupies [Perez 2002]. The relevance is not the historical periodization, which this thesis does not adopt, but one mechanism: financial capital builds the infrastructure during installation, the financing structures outrun the productive economy's ability to absorb them, the Turning Point reprices the installed base, and deployment then proceeds on the repriced substrate. Two consequences for this model. First, the four states and the capital cycle interact rather than nest: a Turning Point arriving during State 2 (physical inflation blocking repair) meets a central bank with less freedom to absorb it than one arriving during State 1, and the

same clearing event lands differently on each. Second, the sequencing discipline the state model already enforces on the reader — *recession first, repression later is a path, not a promise* — is exactly the discipline Perez's cycle adds: the two frameworks make the same demand, that the destination not be mistaken for the next observation. The clock is named here so the phase gates of Section 29.3 can be read against it.

A cross-cutting capacity: duration absorption. The transition from State 1 to State 3 is written above as "repair does not restore a durable marginal buyer," and Section 3.1.3 says what that buyer is made of. Increasingly it is demand commissioned by benchmark rules and delivered by leveraged intermediaries financed in repo. The quantity that decides whether repair holds is therefore duration-absorption capacity, not the debt ratio. The mechanism, and why it binds without warning in the yield, is stated there and not repeated here. Two consequences for the state model follow. First, the observable that distinguishes a State 1 that is holding from a State 1 that is being carried on leverage is the configuration of the plumbing series in Appendix I.5, not the level of the long yield. Second, a policy move taken on the reading of a yield that the plumbing has filtered (Section 13.7) is a State 1 action that can precipitate the State 3 condition it was meant to forestall, by tightening into the mounting. None of this times the transition. It says which series would show it first, and that the yield is not among them.

29.3 Four Phases of Adoption

The curve runs as four overlapping phases, each pulled in by a different constituency. These are **metric-anchored, not calendar-anchored**—transitions happen when telemetry shows they have happened, not when a date arrives:

Phase I: Cypherpunk-led. Plumbing proves out publicly; loop works without permission.
Phase II: Allocator-led. Budgeted workloads close the fee loop; capacity becomes a line item.
Phase III: Composability-led. Receipts become default integration; invisible infrastructure.
Phase IV: Policy-led. Receipts become recognized evidence; jurisdictional differentiation.

These phases overlap and vary by geography and sector. The pattern is robust: early idealists, then necessity-driven budgets, then invisible defaults, and finally institutional and policy recognition. **Calendar horizons are intentionally omitted**—they age badly and invite cheap dunking. What matters is whether the phase gate metrics (below) are met.

29.3.1 Phase Gates and Failure Gates

The following table makes the adoption curve **auditable**. Each phase has entry gates (observable thresholds) and failure gates (conditions that falsify the phase or regress to a prior phase).

Phase Entry Gates:

Phase	Gate Metrics	Threshold
I → II	VERIFYPRICE exists for ≥ 3 canonical workloads with ≥ 6 months history	p95 within SLO
	VERIFYSETTLE for ≥ 1 corridor: protocol-attributable refund failures $= 0$; overall refund_safe $\geq 99.5\%$ with root-cause classification published[b]	success $\geq 95\%$
	VerifyReach published for major regions	$\text{succ}_1 > 0.7$
	Receipt volume non-trivial	≥ 1000 receipts/day
II → III	Fee coverage $\geq 30\%$ of security budget[a]	Sustained 12–24 months
	Budgeted workloads $\geq$ speculative in volume	Visible in receipt tags
	VERIFYSETTLE survives ≥ 1 policy shock	SLOs hold
	Capex/OPEX lines reference verification	Public disclosures exist
	Where agent-originated demand is material: independently originated share rising; HRS and NPCR inside their published bounds (Section 26.8.9)	Principal Map published; bounds pre-declared
III → IV	Collateral/reserve usage in multiple venues	≥ 3 independent
	VERIFYPRICE/Settle stable through macro stress	SLOs hold during crisis
	Duration-neutral holding cohort	Measurable in on-chain data
	Fee coverage $\geq 50\%$	Sustained
IV	Legal codification of receipts	≥ 2 jurisdictions
	Public-sector participation	Visible in dashboards
	Geographical differentiation	Measurable in VerifyReach

Table 29.1: Phase entry gates for adoption curve.

Calibration note on the 30% gate. The II→III fee-coverage threshold sits roughly an order of magnitude above anything achieved by the closest precedents (Filecoin, Render, Akash, Golem, Livepeer), whose native fee revenues have covered low single digits of issuance through their operating lives (Section 34.1.3). The number is not derived from those systems' best outcomes; it is set at the level below which the security budget is issuance-funded and the asset is economically a service token. It should be read as a deliberately hard target with a known failure rate, and Red Line 5 is its enforcement mechanism. [a]*Coverage here is FeeCoverage as defined in Section 25.6.6 — retained fees over budget, burns excluded — not a fees-plus-burn sum; a "fee+burn" sum is not the gate.* [b]*Refund safety is read as Part IV now defines it — a bounded-time, no-loss exit for the honest party — rather than as a single ratio equal to one. A refund failure attributable to the protocol (a timelock, adaptor, or script defect) is disqualifying at any count; failures attributable to the user's own environment (lost keys, a wallet offline past the refund window) are counted, classified, and published, and the corridor passes at* $\geq 99.5\%$ *overall. A gate that demanded exactly* 1.0 *over every failure class would be met by a corridor nobody used. Two denominators are in use and both are published: the Layer 5 SLO and corridor-admission figure (*$\geq 99.9\%$*, Section 23.5) excludes user-environment failures from the denominator; the gate here and Red Line 2 read the* overall *figure (*$\geq 99.5\%$*), which counts them. The stricter-looking number is the narrower denominator, not a stricter test.*

The fee base is a deployment-phase quantity. The 30% gate measures budgeted workloads closing the fee loop — verified compute purchased as a line item by ordinary budgets. On the capital-cycle reading of Section 29.2, that demand is disproportionately a *post*-Turning-Point phenomenon: broad diffusion of a technology into routine procurement follows the repricing event that makes its installed capacity cheap, and precedes nothing. During installation, frontier-lab and hyperscaler demand dominates, and that demand is exactly the bypass channel of Red Line 6 — hyperscaler capacity priced at negotiated enterprise rates the reference panel cannot observe, purchased through channels that never touch the native asset. The phasing claim the thesis should therefore state explicitly, and now does: **Phase III fee coverage is not a smooth path from Phase II; it is a quantity that plausibly becomes measurable only on the far side of a financial clearing event the thesis does not predict and cannot time**. This is a re-timing of the existing gate, not a relaxation of it — the threshold is unchanged, and the regression gate (fee coverage collapses below 30% with workload mix $> 80\%$ speculative) still fires on the same observable. What changes is the expectation: a reader tracking Phase II telemetry through a disorderly macro period should expect the fee base to look *worse* before it looks better, because the demand that feeds it is being built by the same event that is destroying the financing that currently funds it. The trip wire for the opposite reading — that the clearing event came and went and the fee base did not arrive — is Condition B of Red Line 14: real fee turnover failing to grow while physical throughput grows is the deployment phase failing to materialize, measured in public. The implication for anyone planning against this document should be stated as plainly as the dependency: **Phase III is not reachable on a 36-month horizon, and every deliverable that can be dated is a Phase I artifact** — dashboards, corridors, receipt schemas, reference panels — whose value does not depend on the fee base arriving.

Condition	Consequence
VERIFYPRICE p95 exceeds SLO for core workloads for ≥ 3 months with no credible remediation	Red Line 1; SoV thesis falsified
Any protocol-attributable refund failure on an admissible route, or overall refund_safe $< 99.5\%$ (repeated, not isolated)	Red Line 2; Private Money claim fails
Verification concentrates $> 70\%$ on single profile/jurisdiction for ≥ 3 months (reads after the Phase II entry gate; a watch indicator before it)	Red Line 3; SoV thesis falsified
Receipt data becomes unverifiable (datasets unavailable, dashboards dark)	Telemetry capture; thesis unverifiable
Fee coverage falls below 30% and workload mix becomes $> 80\%$ speculative	Phase II $\rightarrow$ I regression
Fee coverage falls below 10% with workload mix $> 80\%$ speculative for ≥ 12 months	Red Line 5; SoV thesis falsified

Table 29.2: Failure gates for adoption phases.

Failure Gates (Phase Regression or Thesis Falsification):

Phase regression logic: Phases can regress. If Phase II metrics collapse (fee coverage falls, budgeted workloads disappear, corridors fail stress tests), the system slides back to Phase I. This is not failure of the technology—it is failure of adoption. The telemetry makes it visible so stakeholders can respond.

29.3.2 Adoption Curve To Expect

The adoption curve for this new monetary substrate will not be a single "flip" moment. It will unfold in discernible phases, each with its own constituencies, failure modes, and telemetry.

The through-line is simple. Under Bretton Woods II, you ran money by buying reserves, hiring lawyers, and trusting gatekeepers. Under Verification, you run money by buying capacity: privacy capacity, proof capacity, verified FLOPs. **Budgets stop paying for promises and start paying for work that anyone can check.**

"Phase I–IV" are not impressions. Each phase has rough, observable thresholds in the same metrics already defined (VERIFYPRICE, VerifyReach, VERIFYSETTLE, FERs, Work Credit utilization). The point of the curve is that you can tell, from dashboards and receipts, which world you are in.

29.3.3 Phase I: Cypherpunk-Led (Prove the Plumbing in Public)

Phase I belongs to the people who already live in the future: open-source cryptographers, hardware hackers, privacy-chain communities, and early AI+ZK builders. Their job is to prove that the loop (Create/Compute → Prove → Settle → Verify) can be made real **without permission**.

In this phase, the main artifacts are reference implementations and receipts:

- **Open-admission prover markets** stand up on top of existing chains and zk networks. They publish live VERIFYPRICE dashboards for the canonical Phase I set {PROOF_2^{20}, MATMUL_4096_FP32_FR — the same set the Phase I trigger below reads on.
- **MatMul-PoUW testnets** (Duplex-style patterns) test whether verification overhead sits inside its modality band (Section 22.2.1) on commodity hardware — $r(W) = v/p \leq 10^{-2}$ for Freivalds-class probabilistic checking [Freivalds 1977] (M2, soundness parameter published), $\leq 10^{-3}$ where a succinct proof is available (M1) — and whether miners without favoured hardware profiles can still win blocks. The legacy figure $r(W) \leq 0.3$ is an engineering floor on which no monetary claim reads; a MATMUL workload passing it and nothing tighter is in the regime where participants sample or delegate rather than check. This is a target under investigation, not a demonstrated result: the optimal-security claim remains conjectural [Komargodski & Weinstein 2025] (Table 22.5), and Phase I exists to produce the measurement rather than cite it.

- **PoL pilots** (Ambient-like patterns [Ambient 2024]) demonstrate hybrid verified inference with honest-output rates that can be measured and contested, not marketed.
- **zk-PoW networks** (Nockchain-style patterns [Nockchain 2025]) act as public receipt ledgers (places where proofs from many domains can be anchored and timestamped).
- **Non-custodial BTC↔XMR swaps with refund-safe UX** move from GitHub curiosities to tools that people actually use. BTC↔XMR is the deployed corridor [COMIT 2021]. A BTC↔shielded-ZEC adaptor-signature swap does not yet exist — shielded Zcash has no adaptor-compatible signing path in production — and that corridor is gated on the cryptographic work, not on UX.
- **Federated e-cash mints** (Fedimint-pattern [Fedimint Project 2024–25]) extend Bitcoin's privacy into community-scale banks: each federation converts Bitcoin into blinded e-cash with full privacy for internal transactions and Lightning for cross-federation settlement. This is *federated custody* — a known federation of guardians holds the bitcoin, and posts no slashable bond — and it belongs in Class B or Class C of the bridge classification (Section 23.4), not in the non-custodial corridor plumbing above. It is listed here because federations disperse the custodian and the blinding hides the internal graph, which is real privacy at community scale. It is not listed as settlement the thesis can count: Red Line 2 rejects custody as the foundation of Private Money, a federation that fails is a custodial failure, and Fedimint e-cash does not count toward any refund-safety or collateral reading. The tension is left standing rather than resolved.

(Note: "Duplex-style," "Ambient-like," and "Nockchain-style" are referenced as design patterns, not endorsements of any specific project or ticker.)

Nothing is "mainstream" yet. Most users are still speculators and hobbyists. But a few things become hard facts rather than hopes: you can buy proofs as a service from permissionless markets; you can pay for them over privacy rails without custody; and anyone with a laptop can verify the receipts.

Phase I telemetry and triggers: **Phase I is real** (not hypothetical) once:

- **VERIFYPRICE exists for a small set of canonical workloads.** At least a handful of public markets publish VERIFYPRICE(W) dashboards for $W \in \{$PROOF_2^{20}, MATMUL_4096_FP32_FREIVALDS, with months of history, and p95 times stay within stated SLOs under stress. (Large-model inference enters this set only when ZKML maturity permits; Section 22.10 gates it. Expecting 70B-class verification on laptop hardware in Phase I would trip Red Line 1 at launch and thereby test nothing but the checklist.)
- **VERIFYSETTLE is measured for at least one serious corridor.** For a BTC↔XMR corridor C — the only non-custodial privacy corridor deployed today; BTC↔shielded-ZEC waits on adaptor-compatible shielded signing — VERIFYSETTLE(C) is public and hits success$(C) \geq 0.95$ with zero protocol-attributable refund failures and refund_safe$(C) \geq 0.995$ overall,

each failure root-caused and classified.

- **VerifyReach is non-degenerate.** At least one verifier network publishes VerifyReach(N, R) for major regions with real multi-path reachability.
- **Receipt volume is non-trivial.** Receipt ledgers anchor a steady flow of PIDL receipts per day (≥ 1000) across multiple workloads.

Once those metrics are visible, the question is no longer "can this exist?" but "can it scale?"

29.3.4 Phase II: Allocator-Led (Proofs and Privacy Become Budget Lines)

Phase II begins when **allocators** (fund managers, corporate treasuries, exchanges, and large web platforms) start to treat privacy and proofs the way they once treated bandwidth: as recurring operating costs, not science projects.

By this point (**after Phase I gate metrics are met**, not by calendar date):

- Proof and compute networks publish open VERIFYPRICE and reliability dashboards with historical data. You can see, month by month, how p95 verification times, costs, and failure rates behave under load.
- The PAL SDK and settlement adapters have been integrated into serious applications: analytics pipelines, custody stacks, compliance tooling, provenance layers for media.
- Privacy corridors (BTC↔XMR, and BTC↔shielded-ZEC if the adaptor-signature work has landed by then) have polished GUIs, reference libraries, and documented latencies. Refund failures are statistical outliers with post-mortems, not routine hazards.

Allocators do what they always do when facing repression and technological change: they **reclassify**. Instead of framing privacy and proof capacity as speculative tokens, they treat them as line items required to keep operating:

- A bank's AI risk model must be run on verifiable compute with audit-friendly receipts, because regulators now ask for them.
- A media platform must attach cryptographic provenance to high-stakes content, because the liability of not doing so is too high.
- A trading venue must use private settlement rails to avoid leaking its entire order book and client graph.

They do not buy these capacities because they have converted to cypherpunk ideology, but because compliance and risk management now require math, not memos. Under Verification, "do nothing" is no longer the conservative option; it is reckless.

Phase II telemetry and triggers: **Phase II is real** once:

- **Fee coverage reaches a meaningful slice of security budget.** For at least one serious PoUW/proof network, fees tied to real workloads cover $\geq 30\%$ of miner/prover revenue over 12–24 months (FeeCoverage convention: retained fees, burns excluded).

- **Budgeted workloads dominate speculative ones in volume.** In receipt analytics, the *number* of proofs purchased by enterprises grows steadily, even if speculative volume remains higher in nominal terms.
- **VERIFYSETTLE stays inside SLOs through at least one policy shock.** For at least one high-volume corridor, VERIFYSETTLE remains within bounds across a visible policy or regulatory event without catastrophic failure.
- **Capex/OPEX lines reference verification capacity explicitly.** At least some institutions treat Work Credits as a line item (verification spend, privacy spend), not as a speculative asset bucket.

At that point, the triad is no longer "crypto" from the allocator's perspective; it is infrastructure they **have** to pay for.

Two allocator paths, only one of which is monetary. "Allocator-led" is ambiguous, and the ambiguity is dangerous. Institutions can arrive in two very different ways, which look similar in a price chart and nothing alike in the telemetry.

Phase II-A — native allocator adoption. Institutions acquire the base asset; stake it or post it as collateral; pay native fees; use private settlement; purchase proofs or verified compute; participate in native liquidity; and hold self-custodied or directly redeemable claims. This is **monetary adoption**, and it moves every series on the Value Capture Board.

Phase II-B — wrapper-led financial exposure. Institutions purchase spot ETFs and treasury-company equity, use futures and options, and hold omnibus custodial claims, gaining price exposure without native use. This is **market adoption**. It may precede monetary adoption, seed it, or permanently substitute for it (Section 5.6); which of the three is an empirical question answered only by VerifyFlow (Section 26.8).

Phase II-C — agent-led native demand. There is a third way demand can arrive, and it looks like neither of the first two on a price chart or on the boards as previously drawn. Delegation contracts issued by software agents (Chapter 7) pay native fees, post native collateral, settle over refund-safe corridors, and hold reserve balances against future collateral calls. On the Value Capture Board this is Phase II-A. On the Native Monetary Buyer Map it is a self-custodied, loss-bearing, countercyclical cohort. Whether it is monetary adoption depends on a question neither board asks: whose principal is issuing the delegations. Fleet-originated demand is Phase II-B by another route, exposure without agency. Independently originated demand with human-held roots is the first constituency on which Red Lines 17 and 18 could read non-null. The Principal Map (Section 26.8.9) is what tells them apart.

Revised Phase II gate. Assets under management, market capitalization, price level, and wrapper launch count **do not** qualify the system for Phase III. Passing requires positive native fee growth; material native burn or retirement; growing collateral lockup; expanding native

settlement; rising receipt volume; healthy VERIFYPRICE, VerifyReach, and VERIFYSETTLE; a Wrapper–Native Growth Gap below a pre-declared threshold; and no uncontrolled custodial concentration. Where agent-originated demand is a material share of native fee turnover, passing additionally requires the Principal Map readings of the gate table: without them a single operator's fleet can pass this gate on every other published series.

False Phase Transition: Wrapper-Led False Positive

Market capitalization and institutional AUM rise rapidly while native use, fee coverage, burns, and collateral remain flat. The ecosystem appears to have advanced to institutional adoption but remains, economically, in the speculative phase. This is the most flattering way for the thesis to fail, which is exactly why it needs a named gate rather than a footnote.

False Phase Transition: Fleet-Led Composability

Receipt volume, native fee share, and collateral lockup all rise because one operator's agent framework routes through the protocol by default; Cross-Perimeter Share is near zero or not computable and the Human-Revocable Share is undefined. The Value Capture Board reads Phase III; the Principal Map reads Phase I with a large customer. Publish the second reading.

Phase III can arrive before Phase II. The four phases are ordered by constituency: cypherpunks, allocators, composability, policy. Software agents do not need a user interface, a treasury committee, or a Phase II narrative; they need a receipt schema, a refund-safe corridor, and a unit anyone can check. Phase III's defining property—receipts as default integration, invisible infrastructure—is what an agent framework ships by construction. Phase III behavior can therefore appear before Phase II's allocators do. The gates are metric-anchored precisely so that ordering does not matter, no gate is relaxed by this observation, and the false-positive test is the same as for Phase II-B: composability that is one operator's SDK is a fleet, not a phase.

29.3.5 Phase III: Composability-Led (The Stack Disappears into Infrastructure)

Once Phase II gate metrics are met and sustained, composability becomes second nature. Developers no longer think in terms of "ZK project X" or "PoUW chain Y"—they think in terms of the loop.

Provenance, verified compute, and private settlement start to resemble TLS on the web:

- New applications default to emitting claims and receiving receipts via the SDK, simply because it is easier than hand-rolling trust.
- Major frameworks and toolchains ship with verification modules and privacy adapters bundled: verifying a receipt feels as ordinary as opening an HTTPS socket.

- Chains and rollups routinely outsource heavy computation to PoUW or proof factories and anchor receipts on shared ledgers.

From the outside, nothing dramatic happens. There is no "flippening." What changes is the **default**:

- Sensitive data is processed either under enclaves with open silicon profiles or under ZK, and accompanied by attestations.
- Payments that cross borders or touch politically exposed persons quietly route over privacy rails, with receipts that satisfy auditors but not censors.
- AI systems whose outputs matter are either run with PoL-style verifiable inference or backed by helmets of audits and shadow runs, with failure rates visible rather than buried.

Networks that can demonstrate a clear, measurable link between work and value, paired with credible scarcity and decentralization telemetry, are the ones that establish **durable cash-flow accrual** to the asset. That is what this phase can show. Whether a monetary premium sits on top of it is a separate question, decided by the holder-side flow of Section 13.3.11 rather than by the accrual, and not one that price alone answers in either direction (Corollary 13.1).

Phase III telemetry and triggers: **Phase III is real** once:

- **Non-trivial native collateral and reserve usage.** The base asset appears as collateral or reserve holdings in ≥ 3 independent native venues, while Work Credits remain separately reported service claims. The Native Monetary Buyer Map shows a visible self-custodied, loss-bearing cohort rather than inferring reserve status from lockup.
- **Macro sensitivity flips.** Price and flows react as much to changes in workload budgets (AI capex, compliance requirements, settlement volume) as to crypto-native news.
- **VerifyPrice and VerifySettle stability through macro stress.** During interest-rate shocks or liquidity crunches, the underlying utility, capacity, and safety metrics do not collapse.
- **Duration-neutral base-asset holding behavior.** Holder data shows a meaningful self-custodied cohort holding the base asset for agency and SoV properties, with holding periods measured in years; Work Credit inventory is reported separately as service demand.
- **Fee coverage $\geq 50\%$.** Sustained over multiple quarters.

When these show up in the dashboards and market structure, the base asset has advanced as a monetary candidate. The triad has demonstrated service relevance; Work Credits have not changed instrument class.

29.3.6 Phase IV: Policy-Led and Path-Dependent (Verification as Public Good)

Once Phase III gate metrics are sustained through macro stress, the curve becomes more path-dependent and political. If the stack delivers on its risk promises (decent decentralization, cheap verification, robust privacy corridors that are visibly abuse-resistant), states and institu-

tions will begin to treat parts of it as **public goods** rather than threats.

Some jurisdictions will move first. They will:

- Codify cryptographic receipts as acceptable evidence in courts and regulatory filings.
- Mandate provenance proofs for certain classes of media or AI systems rather than ad hoc labelling.
- Recognize privacy rails with viewing-key regimes as compliant infrastructure rather than as dark pools.

Others will resist, preferring the comfort of chokepoints and legacy gatekeepers. That is where the anti-repression design matters. If privacy rails remain non-custodial and censorship-resistant, if proof and compute markets remain globally addressable, capital and talent can route around laggard jurisdictions the way data routed around telcos.

In the best case, the monetary role of the triad becomes self-reinforcing:

- Each wave of repression (negative real yields, capital controls, information crackdowns) pushes more savings and more workflows onto rails where verification is cheap and permission is irrelevant.
- Each wave of adoption increases the depth and liquidity of proof and compute markets, which in turn makes it cheaper and more obvious to use them for new domains.

If the triad earns a store-of-value premium in this scenario, it is not because everyone suddenly shares a philosophical vision, but because the world has learned, through trial and error, that capacities which are verifiable, censorship-resistant, and continually demanded are safer long-term anchors than promises that can be administered or revoked. This paragraph describes a best case, not a forecast, and the premium in it remains the unsized quantity of Section 13.3.11.

Phase IV telemetry and triggers: **Phase IV is real** once:

- **Legal codification of receipts and corridors.** Multiple jurisdictions explicitly recognize PIDL-style receipts as valid evidence in regulation and courts; some codify acceptable VERIFYSETTLE or provenance requirements into statute.
- **Regulatory reliance on cryptographic proofs.** Policies and enforcement actions reference proof primitives (ZK proofs, PoL audits, receipt formats) rather than merely "records" or "reports."
- **Public-sector participation.** Public institutions (development banks, treasuries, public broadcasters) use triad-aligned rails and publish their own receipt and corridor metrics.
- **Geographical differentiation.** A measurable share of volume and capital migrates toward jurisdictions and networks that respect lawful privacy and verifiable compute, as seen in regional splits in VerifyReach, settlement volume, and Work Credit usage.

This adoption curve is not guaranteed. Verification cost can creep; hardware or router oligopolies can re-centralize control; regulation can pinch on- and off-ramps harder than ex-

pected. But if the stack is anchored in verifiable machines, keeps verification cheap and public, treats privacy as infrastructure rather than vice, and forces its own claims through telemetry and receipts, this is the curve it can plausibly aim at.

29.3.7 The Negative Case: Useful Infrastructure Without Money

It is important to describe what it looks like if the triad remains useful infrastructure but never becomes money:

- Proof markets exist and grow, but fees are paid in stablecoins or fiat. The native asset has weak demand.
- Privacy corridors work, but most users route through custodial services that accept traditional payment.
- Verified compute is valued, but hyperscalers dominate and price in USD. Decentralized alternatives remain niche.
- The base token trades as a speculative asset correlated with crypto markets, not as a store of value with independent demand.
- Fee coverage never reaches 30%; the security budget depends on issuance subsidies indefinitely.

In this scenario, the *stack* succeeds (it delivers real privacy, proofs, and compute) but the *monetary thesis* fails (the native asset does not capture the economics). This is the "utility-token trap" outcome. It is not catastrophic—the infrastructure is still valuable—but it means the store-of-value claim was wrong.

The telemetry makes this outcome **visible**: fee coverage, native-asset fee share, bypass-channel indicators, and holder behavior all signal whether the monetary thesis is tracking or not.

This scenario is not hypothetical. It maps directly to one pole of a live debate among Bitcoin practitioners: the position that Bitcoin succeeds as an escape hatch and self-custody tool for a small elite, but does not disrupt the existing power structure or impose a free market. The thesis's value-capture conditions and telemetry regime exist precisely to make this distinction measurable rather than a matter of temperament.

29.4 Failure Gates

Phases can regress. If Phase II metrics collapse, the system slides back to Phase I. The failure-gate conditions and consequences are stated once, with the phase entry gates in Section 29.3; they are not repeated here. The regression logic is worth restating because it is the part that gets forgotten: regression is not failure of the technology—it is failure of adoption, and the telemetry exists so that stakeholders can see the difference.

Chapter 30

Risk Analysis & Failure Modes

The risks to this emerging triad are not merely technical; they are structural, political, and economic.

30.1 Technical Risks

Every guarantee in this thesis reduces, at the bottom, to mathematics running on matter that somebody manufactured. That is where the first class of risk lives — not in the cryptography, which is either sound or it is not, but in the supply chain underneath it.

Proof system failures. A SNARK/STARK is broken, or a circuit or verifier fails to constrain what it claims: a missing constraint in Zcash's Orchard circuit could have allowed counterfeiting that cannot be detected after the fact [Zcash 2026]. The blast radius is bounded before it happens: multi-ZK support, proofs tagged with system and parameter IDs so that a failed system can be delisted, a turnstile on each in-asset shielded pool that caps what can leave it at the value that entered, and migration paths mean one failure deprecates one proof family or circuit rather than the entire attestation layer.

Hardware capture. A major fab or TEE platform ships a backdoor. This is the Layer 0 argument wearing different clothes: profiles, lot sampling, and open hardware alternatives exist so that one compromised supplier cannot silently become the substrate of record for every receipt (Section 17.4).

Protocol bugs. Consensus bugs, bridge flaws, privacy leaks. Mitigation: reference implementations, staged rollouts, incident response.

The common thread is that none of these risks is existential on its own; each is existential only if it arrives *unobserved*. A broken proof system caught in staging is an engineering incident. The same break discovered two years after deployment, with every receipt since built on it, is a monetary event.

30.2 Economic Risks

The economic risks are quieter than the technical ones, and more dangerous for it. They do not announce themselves as failures; they arrive as slow drift that every individual reading can explain away.

Verification cost creep. If verifying proofs ceases to be much cheaper than producing them, the entire asymmetry collapses — and it collapses gradually, one mildly disappointing benchmark at a time. Shows up as rising $r(W)$ in VERIFYPRICE; the trip wire exists because no single reading looks fatal.

Centralization. When specialized hardware or closed routers dominate, neutrality erodes. Shows up in concentrated facility IDs, rising entry latency.

Fee death-spirals. If demand falls, fees collapse, security budgets shrink. Mitigation: conservative base issuance, Work Credits encoding long-term demand.

Paying the guards in fresh scrip. The fee-coverage question deserves the concrete form, because the table of 30/50/80% targets is easy to read as a KPI and forget. A system that relies entirely on issuance to fund security is a warehouse paying its guards each morning in shares the warehouse prints that morning. The guards keep showing up — right up until the shares stop buying lunch. The fee-coverage trajectory is simply the schedule for switching them to wages.

30.3 Political Risks

Political risk is the one class this thesis cannot engineer around, only prepare for. The adversary here is not a bug or a market; it is an institution with jurisdiction, patience, and every incentive to frame the stack in the least flattering available light.

Bans & sanctions. Jurisdictions may declare privacy assets illegal, sanction contracts, or criminalize the use of certain clients. Mitigation: jurisdictional diversity (labs, foundations, hardware profiles, LPs); client modes that degrade gracefully (offline, mesh, sat-links); legal defense resources; clear separation between core protocol and specific front-ends so that UI bans do not equal protocol death.

Info-ops and framing. Media and governments can frame the stack as "criminal tech" or "national security threat." Mitigation: lawful-privacy narratives grounded in receipts and compliance primitives; visible legitimate use cases (payroll, provenance, AI verification, critical-infrastructure audit trails); and an insistence on evidence ("show the receipts") rather than slogans.

Regulatory chokepoints at the edge. App stores, banks, and ISPs can be pressured to block access even when the protocol is neutral. Mitigation: the Layer-1/2 work described earlier (al-

ternate transports, side-loading, content-addressed and offline distribution) so that no single storefront, bank, or carrier can become a kill switch.

Political risk cannot be "engineered away." It must be **distributed and prepared for**: diversified jurisdictions, many independent implementations, multiple access paths, and a culture that expects (and drills) for attempted bans and smear campaigns rather than treating them as unthinkable. The stack is built so that being declared unwelcome in one jurisdiction is survivable; no layer assumes it will never be declared unwelcome anywhere.

30.4 Spec Drift, Vaporware, and Scoreboard Capture

Spec drift and vaporware can hollow out trust long before censorship does. Whitepapers without shipped code or inflated performance claims corrode the ecosystem's credibility. In metrics, this shows up as:

- a widening gap between claimed and measured $\textsc{VerifyPrice}(W)$;
- a mismatch between Work Credit issuance and actual workload utilization;
- chains with high market caps and low receipt volume;
- clients that are nominally "live" but barely process real workloads.

Active liveness tracking (clients, explorers, throughput under real load, settlement safety metrics) should replace social metrics or TVL as the industry's scoreboard. Projects that publicly document testnet-to-mainnet milestones and publish SLOs and incident reports exemplify the discipline needed.

No dashboards, no trust. If $\textsc{VerifyPrice}$, swap success and refund safety, decentralization telemetry, and corridor health are not public, treat claims as unpriced risk. A system that refuses to show its receipts is asking you to underwrite the very soft guarantees this thesis is designed to escape.

30.4.1 Scoreboard Capture in the Market Realization Plane

Scoreboard capture has a financial-product form as well as a protocol form, and it is the more seductive of the two. A product industry under competitive pressure increasingly asks "what ticker might catch a flow?" rather than "what does the investor actually need," with capital chasing realized performance and sponsors launching ever more specialized wrappers. The result is a scoreboard on which a wrapper can be a triumph while its holders are destroyed — reverse splits keep the share price presentable, new inflows replenish assets, and fees keep accruing to the sponsor (Section 26.8.4).

Scoreboard Capture, Extended

Scoreboard capture occurs when product AUM, market capitalization, launch count, or trailing returns substitute for measurement of whether *capital survived*, *native services were used*, and *the monetary loop functioned*.

The countermeasures are the Capital Survival Ratio, the Wrapper–Native Growth Gap, and Native Use Share (Section 26.8). None of them can be gamed by launching another product.

30.5 Collateral Architecture Under Wrong-Way Risk

Collateral design chooses a distribution of failure; it does not eliminate one. At least three models must be reported and stress-tested against the same drawdown and DVC scenarios:

Pure-native collateral. Operators post only the base asset. This maximizes direct lockup and value capture but also maximizes wrong-way risk: price decline raises required units, liquidation or exit releases float, and falling capacity can weaken the service case that supported the asset.

Mixed collateral. Operators post a published mixture of native and external reserve collateral. Direct native capture is weaker, basis and custody risks enter through the external leg, but collateral coverage can remain more stable through a native drawdown.

Native plus insurance. Native stake supplies incentive alignment while a separately capitalized first-loss, insurance, or resolution pool absorbs defined tail losses. The pool must publish capital, exclusions, correlation, claims priority, and replenishment rules; an unfunded promise is not insurance.

The comparison must report collateral coverage against realized slashing exposure, liquidation sensitivity, operator survival, released float, DVC, and service-SLO recovery. Governance may choose among the models only after publishing those trade-offs. It may not describe the strongest value-capture design as the safest prudential design.

30.6 Red Lines: When the SoV Thesis Fails

The thesis is falsifiable. If any of these conditions persist without credible remediation, the store-of-value claim is **no longer defensible**. Severity labels are used consistently below and mean specific things: **kills** = the condition falsifies the thesis outright (the chain's hinge is severed); **weakens** = the condition removes one leg of support such that the thesis survives only in reduced form and must be restated. The label appears in each red line's closing paragraph and is not a tone choice.

Red Line 1: Verification Affordability Breaks

Condition: VERIFYPRICE p95 (the Real-Resource VerifyPrice SLO series, not the Physical VerifyPrice audit-cost series of Red Line 10) exceeds SLO bounds for core workloads for ≥ 3 consecutive months, with no credible remediation path.
Why it kills the thesis: "Anyone can verify" is the hinge. If verification becomes expensive, proofs become platform claims, not public facts.

The touchstone is the image to keep. Gold earned its monetary career on a dark stone against which a nugget could be streaked in seconds — expensive to produce, cheap to assay approximately (touchstone and density; a tungsten core defeats the first and is caught by the second), and therefore expensive to fake against a buyer who runs both checks. The entire premium of this thesis rests on the assay staying cheap. The moment verification costs real money or real expertise, receipts revert to what platform labels already are: assertions with better typography. Three months of sustained breach is not a performance regression; it is the anchor reverting to the broadcast era.

Corollary 30.1 (Delegation Topology)**.** *Red Line 1 reads on whether anyone* can *verify. The same series governs which delegations cross a trust perimeter and which of those buy a proof (Section 7.3). Let $r(W) = v(W)/p(W)$ as in Appendix A. Let $\ell(W)$ be the expected loss from accepting an unverified artifact of workload W, and $a(W)$ the cost advantage of the best outside operator over the parent's own perimeter, both as fractions of $p(W)$. A parent that sends a hop outside either checks the artifact or bears the expected loss, so crossing costs it $\min\{r(W), \ell(W)\}$. It crosses when $a(W) > \min\{r(W), \ell(W)\}$, and it buys verification when it crosses with $r(W) < \ell(W)$. Proof demand from delegation is therefore the set of hops with $r(W) < \min\{\ell(W), a(W)\}$: work that is worth sending out and too consequential to accept unchecked.*

The corollary is a mechanism, not a measurement: $\ell(W)$ and $a(W)$ are workload- and principal-specific and are nowhere estimated in this document. It states a necessary condition and not a sufficient one; a hop that clears it can still stay inside a fleet for want of settlement, reach, or legal cover. Two readings follow. Low-stakes work, with $\ell(W)$ small, crosses perimeters unchecked whenever it is cheaper outside, and buys no proofs at any $r(W)$: not every delegation is a customer. And for consequential work, a rise in $r(W)$ removes hops from the open economy in order of their outside advantage, long before verification costs as much as production. A VERIFYPRICE regression that stays inside the constitutional SLO but crosses $a(W)$ for the workloads carrying most consequential hop volume has not tripped Red Line 1 and has already produced enclosure. The Cross-Perimeter Share (Section 7.2) is the series that would show it: falling CPS with VERIFYPRICE inside SLO is the signature of enclosure by convenience. No threshold of Red Line 1 moves, and the corollary trips nothing on its own.

Reading order. Red Line 11 is downstream of Red Line 1 in a way the two lines do not otherwise state. Capacity concentration is the lagging indicator; falling CPS is the leading one; VERIFYPRICE drift toward $a(W)$ is one cause, and a closed stack shipping enough capability to shrink $a(W)$ is the other. Once CPS is computable, the Layer 0 Capacity Board should carry the three series in that order.

What would falsify the corollary. Rising CPS on consequential workloads while $r(W)$ is stable or rising for the same workloads and outside prices have not fallen relative to inside ones, sustained over four quarters. That would mean delegation is crossing perimeters for reasons other than verifiability—contract law, reputation, insurance—and that the claim in Section 7.2, that verification is the only reliability mechanism that fits inside the hop, is wrong for those workloads. Where the price advantage cannot be separated from the verification cost the reading is inconclusive and is published as such. The thesis should welcome the falsifying result: it would mean the institutions Part I says are weakening are still doing the work.

Red Line 2: Refund Safety Breach

Condition: Any protocol-attributable refund failure (a timelock, adaptor, or script defect that leaves an honest party without a bounded-time, no-loss exit) on an admissible route, or overall corridor refund_safe $< 99.5\%$ with failures unclassified, with repeated incidents and no automatic delist + remediation. Refund safety is read as Section 23.4 defines it; the same specification gates Phase I in Section 29.3.
Why it kills the thesis: Non-custodial settlement is the "Private Money" foundation.

One lost deposit is an incident. A corridor that eats funds and remains listed is a verdict: the system has told every user what it actually is. On the protocol's own conduct refund safety is binary by design — there is no partial custody, no 97% non-custodial — because trust does not survive being mostly honored; the tolerance in the overall figure exists only for failures the user's own environment caused, each of which must be root-caused and published to count as such. The automatic delist matters more than the breach itself; a network that keeps a broken corridor on its board for volume's sake has chosen the scoreboard over the user.

Red Line 3: Verification Monoculture

Condition: $> 70\%$ of verifications on a single hardware profile, TEE vendor, or jurisdiction for ≥ 3 consecutive months.
Maturity gate. This line reads only after the Phase II entry gate of Section 29.3 is met; before that it is published as a watch indicator and trips nothing. A Phase I network with three verifier operators in one country is concentrated by arithmetic, not by capture, and a line that fires at launch tests the checklist rather than the thesis.
Second reading. A migration of the verified-capacity mix from M1 toward M3 ends at

the same place — "trust the verifier" — without any single vendor dominating. That reading is instrumented separately, as the M1+M2 share of verified units under Invariant 1(a) (Section 25.4); a fall below its published floor is Sev-1, and the two readings are published side by side.
Why it kills the thesis: Monoculture means "trust the dominant vendor."

Seven-tenths of the world's verifying running on vendor-attested chips from one supplier is not an efficiency; it is a single point of political failure wearing a supply-chain dashboard. The concentration can arrive without any malicious act at all — one vendor simply being better for two quarters is enough — which is precisely why the line is drawn on the reading, not on intent. The moment verification monocultures, the adversary no longer needs to break the cryptography; it needs only to lean on the vendor. The threshold here is deliberately looser than Red Line 11's 50% on verified compute capacity, and the difference is the two-exposures distinction of Section 17.10.1: verification runs on small quantities of reference hardware at the edge and can be relocated in weeks, so concentration in it is reversible and the line tolerates more of it before reading capture; proving capacity is capital-intensive, interconnection-bound, and slow to move, so the same concentration there is a structural fact rather than a passing one and the line reads earlier.

Red Line 4: Telemetry Capture

Condition: Receipt datasets become unavailable, unverifiable, or controlled by a single party.
Why it kills the thesis: If telemetry can be captured, the entire observability regime is theater.

A constitution nobody can check is a press release. Every instrument in Part V exists on the assumption that an independent party can re-derive the published number from raw public artifacts and get the same answer; the moment the raw data itself becomes a proprietary good, the dashboards continue to render but stop meaning anything. This is the quietest of the kill conditions — no price moves, no users are harmed, nothing visibly breaks — and it is the precondition for every other red line's measurement. Capture the scoreboard and every other game on it is rigged.

Red Line 5: Fee Coverage Collapse

Condition: FeeCoverage (retained fees as a share of security budget; burns excluded per Section 25.6.6) falls below 10% and workload mix becomes $> 80\%$ speculative for ≥ 12 months. The 30% figure in Chapter 29 is a phase-regression gate, not this red line.
Why it kills the thesis: The SoV story requires structural demand, not pure narrative.

This is the guards-in-fresh-scrip condition (see Section 30.2) promoted from metaphor to trip wire. A network can coast on issuance for years, and during those years every headline number looks healthy: hash rate, market cap, even transaction counts. What it cannot do is fake the composition of its demand. When four out of five units of workload exist to farm a token rather than to be used, the fee floor is not being tested — it is being performed, and the security budget is a subsidy awaiting its first serious drawdown.

Red Line 6: Value Capture Failure

Condition: Triad usage grows (proof volume, settlement volume, compute demand) but native-asset fee volume, burns, collateral lockups, and fee coverage do not grow with it for ≥ 12 months.
Why it kills the thesis: The system may be useful infrastructure but not a store-of-value asset. Users are consuming triad capacity through bypass channels.

The paper-gold precedent is the cautionary case, and it is a century old: under the interwar gold-exchange standard, and later in the London bullion market's unallocated accounts, claims on gold outgrew holding of gold, intermediaries multiplied paper while vaults stood still, and the price of the paper discovered it could move without the metal underneath it. The same substitution is available here, and the telemetry exists to watch it arrive. Usage growing while native fee volume stagnates is not a paradox — it is the market routing around the asset while continuing to use the network, and it is the single clearest signal that the thesis's monetary link has quietly severed while the infrastructure succeeded.

Red Line 7: Legal Incompatibility

Condition: Lawful users in major jurisdictions cannot use privacy rails without unacceptable compliance uncertainty for ≥ 12 months, and no lawful-privacy patterns (viewing keys, scoped disclosure) are adopted.
Why it weakens the thesis: If regulated actors cannot participate, the "Private Money" leg loses its institutional constituency and anonymity sets shrink.

Privacy that only outlaws can use is a shrinking set, not a monetary property. The thesis's lawful-privacy design exists precisely to avoid the corner where compliance uncertainty meets default privacy and the only remaining users are the ones who were never complying. Anonymity is a network good: each legitimate departure thins the crowd that makes blending plausible, and the set collapses gradually — which is why this line weakens rather than kills, and why its telemetry (corridor compliance receipts, viewing-key usage) reads on adoption, not on law.

Red Line 8: Governance Capture

Condition: Governance can alter issuance schedules, fee routing, or telemetry rules without hard constitutional constraints, timelocks, or supermajority requirements; *or* a single address family holds more than 20% of time-locked governance weight (Section 25.5) for $\geq$ 2 consecutive quarters — the concentration threshold at which the per-family cap on override votes binds; *or* the narrow emergency path (Section 25.5.3) becomes routine rather than exceptional — more than two invocations in any four consecutive quarters, or any invocation whose subtractive-only scope is eroded to permit value redirection. The structural clauses are binary; the two numeric clauses are what the boards read.

Why it kills the thesis: If insiders can redirect value away from holders or inflate supply at will, the asset is a platform token, not a store of value.

The constitution here is whatever governance cannot silently violate without getting caught (Section 25.1 states the principle). A body that can only shrink the protocol cannot steal from it — that is the entire logic of the subtractive-only emergency path, and the reason its erosion is a kill condition rather than a governance hygiene note. The failure arrives as precedent, not as heist: one emergency invoked for convenience, one scope quietly widened, and the constraint that made the asset safe to hold has been converted into a knob.

Red Line 9: Wrapper Dominance Becomes Monetary Substitution

Condition: Wrapper Dominance Ratio (Section 13.6) rises for $\geq$ 2 consecutive quarters while native fee share, private settlement volume, and collateral lockups stagnate or decline. The WDR is a *pair* — its stock component (custodial-exposure share of total exposure) and its flow component (wrapper-activity ratio) — and "rises" means *both components* rise over the window; if one rises while the other falls, the condition is not met and the split is published. In its full form: custodial and synthetic exposure (CCR, SER) becomes the dominant form of ownership for a sustained period while protocol-native fees, collateral, private settlement, proof demand, and verified-compute usage remain stagnant or declining, with a persistently positive Wrapper–Native Growth Gap in its unit-denominated form (Section 26.8) — the USD form is published alongside but is not the trip wire, because it can move on revaluation alone.

Maturity gate. This line reads only after the Phase II entry gate of Section 29.3 is met; before that it is published as a watch indicator and trips nothing. A wrapper that is the only legally holdable form of a Phase I asset dominates a native base that does not yet exist, and the ratio says nothing until there is a native loop for it to substitute for.

Why it weakens the thesis: Price rises while the monetary thesis dies. Government-approved wrappers becoming the dominant institutional form is co-option without cen-

sorship. The asset may remain a successful financial product, but the claim that it functions as a native store of value *for Privacy, Proofs, and Compute* is no longer supported by anything except its chart.

Red Line 10: Physical Infrastructure Opacity

Condition: Physical VerifyPrice (Section 22.8) — the capacity-weighted p95 time or cost to verify FCR claims — exceeds $2\times$ its benign-state baseline for ≥ 2 consecutive quarters, where the baseline is the trailing four-quarter median published before the window opens and frozen for its duration; or Facility Capacity Receipt data is unavailable or unverifiable for $> 20\%$ of active capacity for ≥ 2 consecutive quarters. Below the Phase II entry gate of Section 29.3 the reading is published as a watch indicator and the baseline is accumulated; the line trips nothing until a baseline exists.
Why it kills the thesis: Layer 0 claims become trust-me claims; Work Credits lose infrastructure credibility.

Red Line 11: AI Enclosure

Condition: Top-3 hyperscalers, closed TEEs, or a single jurisdiction exceed 50% of VERIFYPRICE-tracked verified compute capacity (the 40% chip-family/TEE cap of Section 17.4 being the early-warning tier of the same concentration measure), i.e., the Enclosure Risk Flag (Section 5.6) triggers and persists for ≥ 2 consecutive quarters.
Maturity gate. This line reads only after the Phase II entry gate of Section 29.3 is met; before that it is published as a watch indicator and trips nothing. The 50% threshold is tighter than Red Line 3's 70% for the reason given there: capacity concentration is slow to reverse, verification concentration is not.
Why it kills the thesis: The verified-compute service market described by the "AI Money" lens becomes a cloud IOU; the frontier is enclosed rather than homesteaded, so it cannot support the associated base-asset monetary candidate.

Red Line 12: Agency Failure

Condition: Forced Disclosure Incidence (Section 27.1.2) becomes systemic — viewing keys or full-history disclosure required in $\geq 10\%$ of surveyed employment, banking, housing, travel, or benefits onboarding events in any major jurisdiction, published by the Agency Preservation Board — or the non-custodial share of settled native volume falls below 50% while institutional usage grows, either for ≥ 12 months.
Watch clause (standing). Published beside the two conditions above, and tripping nothing until its threshold is fixed and its instrument is computable: the Human-Revocable Share of agent-controlled native balances (Section 26.8.9) below [50%] for ≥ 4 consecu-

tive quarters *while* the agent-originated share of native fee turnover rises. The bracketed figure is a placeholder; the conjunction with rising agent share is what separates standing inversion (Section 5.2) from a stack nobody uses. The clause reads only after the Phase II entry gate of Section 29.3. It adds a reading to this line and removes nothing from it.

Why it kills the thesis: Lawful privacy collapses socially, or the protocol serves institutions but not users—violating the ninth SoV requirement (Item 9). Emergency governance becoming routine is a related failure: neutrality decays into foundation fiat.

Red Line 13: Energy Sovereignty Failure

Condition: Network capacity-weighted sovereign optionality $\mathcal{O}_s^{\text{net}}$ (Section 17.8) falls below the floor of the published band schedule — the lower edge of the "low" band that carries the issuance cap in Section 25.6.5 — for ≥ 2 consecutive quarters, **or** $> 50\%$ of VERIFYPRICE-tracked capacity sits in jurisdictions operating an active curtailment or power-rationing regime against verification workloads for ≥ 2 consecutive quarters.

Maturity gate. This line reads only after the Phase II entry gate of Section 29.3 is met; before that it is published as a watch indicator and trips nothing. A Phase I fleet is small enough that a single facility's jurisdiction is the network's, and the reading would describe the fleet's size rather than its exposure.

Why it kills the thesis: Verification affordability becomes a sovereign policy variable rather than a market outcome. Red Line 1 becomes *externally triggerable*: an adversary or host state can break the hinge by adjusting tariffs, interconnection queues, or load priority, without touching the cryptography, the governance, or the market.

Red Line 13 and its neighbours. Three red lines now touch the physical substrate and they are not redundant:

- **Red Line 10** is *opacity*: the substrate cannot be seen.
- **Red Line 13** is *fragility*: it can be seen clearly and it will not survive pressure.
- **Red Line 11** is *enclosure*: it can be seen, it is robust, and someone else owns it.

A network can breach any one without the others. Perfect FCR disclosure of a single-interconnect fleet in a rationing jurisdiction breaches 13 while passing 10. Do not merge them.

The dependency this makes explicit. Section 22.7 describes the Real-Resource VerifyPrice SLOs as constitutional and exogenous to token price. That is right, and Red Line 1 correctly treats their sustained breach as fatal. But exogenous to token price is not exogenous to everything: the SLOs remain endogenous to the physical and jurisdictional conditions under which power and reference hardware are obtainable (Section 5.5).

Red Line 13 exists so that this dependency is monitored upstream rather than discovered downstream. Without it, the sequence

$$\text{energy policy} \;\rightarrow\; \text{proving cost and siting} \;\rightarrow\; \text{Real-Resource VERIFYPRICE} \;\rightarrow\; \text{Red Line 1}$$

runs to completion with no instrument reading anywhere along it until the hinge itself fails. A falsification framework whose primary condition can be tripped by an unmonitored external variable is not yet falsifiable in the way it claims.

Red Line 14: The Capturable Wedge Closes

Measured quantities. For canonical workload W at tier T in period t:

- $P^{\text{prot}}_{W,T,t}$ — all-in USD price to the buyer of executing W at tier T through the protocol, inclusive of the protocol fee. Protocol-native: PIDL receipts already carry workload ID, tier, timestamps, and hardware profile (Section 22.3).
- $P^{\text{byp}}_{W,t}$ — lowest publicly quoted USD price to execute the same computational content *without* protocol-grade verification, on the same hardware class. Sourced from a **pre-committed reference panel**: a frozen list of venues (e.g., three hyperscalers and three GPU marketplaces) *plus a fork row — the lowest-fee protocol-grade fork offering the same canonical workload, priced on identical terms*. The fork row is not optional: a fork runs the same technology, so it carries none of the cost handicap a hyperscaler does, and it is therefore the channel that binds first (Section 13.3.8). A wedge measured only against unverified alternatives is measured against the weaker competitor. The panel is scraped daily, with the workload-to-instance mapping published as a machine-readable spec. Panel changes require publication and 90 days' notice. Where a venue publishes committed-use, reserved, or spot pricing, the *lowest generally available* of those is the reading, not the on-demand list price.
- $s = (P^{\text{prot}} - P^{\text{byp}})/P^{\text{byp}}$ — **normalized** bypass spread. Normalization is not cosmetic: the absolute spread shrinks mechanically with compute deflation, so only the ratio is economically meaningful.
- γ — verification cost share: proving plus redundancy per unit, expressed as a fraction of P^{byp}. Derived from FER and VERIFYPRICE telemetry, and subject to the same third-party custody and reproducibility requirement as the reference panel (below), because γ enters ω negatively and the fee recipient has a direct interest in understating it.
- $\omega = s - \gamma$ — the **capturable wedge**: what is left of the price the protocol commands after paying for the verification that justifies it.
- τ — **realized take rate**: *gross* native fees (USD) / (settled protocol turnover, USD). Gross, not net-of-burn, because τ measures what buyers pay — the same convention

the FeeCoverage definition pins for the fee family (Section 25.6.6); burns are supply arithmetic, read as reduced net issuance, and adding them would count destruction as payment. Computable by anyone from chain data.

Condition A — the wedge closes. Capacity-weighted ω across the canonical starter set (Section 22.10) falls below the realized take rate, $\omega < \tau$, for ≥ 2 consecutive quarters, with no credible remediation path.

Condition B — the base deflates faster than volume compensates. Native fee turnover fails to grow over a trailing four-quarter window while protocol physical throughput (verified units delivered) grows over the same window. "Fails to grow" means the trailing-four-quarter sum is non-increasing ($\leq$ its prior value) for two consecutive quarters. Deflation adjustment is by the *workload unit-price index* — the capacity-weighted average USD price of the canonical workload basket — not by the asset price: the condition asks whether buyers are transacting less even as physical delivery grows, so the deflator strips compute-price deflation, not asset-price revaluation. Both the raw and deflated series are published. **Persistence:** the red line reads on the condition holding for two consecutive quarters, not one; a single flat quarter is a flag, not a trip.

The connective, stated once: either condition breaching trips the red line. A breach of A, or a breach of B, or both, is a breach of Red Line 14. There is no requirement that they breach together.

Why A kills the thesis: The protocol is charging more than the differential value it supplies. It is living on switching costs, subsidy, or inertia rather than on anything a buyer would pay for, and the fee is not a wedge on genuine differential value but a tax on captive volume. That fee is retractable by competition, so nothing durable accrues.

Why B kills the thesis: The wedge may be intact and the burn still cannot scale. An ad valorem fee on a deflating unit price is a shrinking real toll (Section 13.3.2); if throughput grows while real gross native fee turnover does not, the empirical bet the thesis is making on volume outrunning deflation is losing, in public.

Both clauses must be monitored, because each closes a gaming route the other leaves open. A published alone is gameable by cutting the fee toward zero, which trivially restores $\omega > \tau$ while capturing nothing; B published alone is gameable by inflating volume at a vanishing wedge. Monitoring both means a protocol that is neither differentially valuable nor growing its real take has nowhere to stand. This is a requirement on the *test*, not on the trigger: dropping either clause from the instrument panel is a breach of the falsification protocol, while breaching either clause in the readings is a breach of the red line.

Why Condition A is hard to game. It is **self-normalizing**. It compares two independently measured quantities — what the market will bear net of verification cost, and what the protocol actually takes — rather than testing either against a threshold somebody had to guess. There is no parameter to lobby for. A protocol cannot pass by lowering its ambitions, because lowering the take rate lowers τ and lowering the price lowers s, and the comparison survives both.

What Condition A is asking, in plain terms. Stripped of the notation, $\omega < \tau$ asks a question with no technical content in it at all: **does anybody actually pay for the truth?** The wedge s is the premium the market will bear for a verified unit over an unverified one, net of what the verification costs. It is a direct measurement of whether cheap checkability changes any buyer's behavior—and it can go to zero for a reason that has nothing whatever to do with this thesis's engineering. Verification can be cheap, correct, public, and comprehensively ignored. An auditable system that nobody audits, and whose auditability moves no price, is transparency theater (Section 13.3.10): the receipts are produced, the proofs verify, every SLO stays green, and no counterparty pays a cent more for the position that carries them.

Why that failure deserves naming rather than only measuring. It has no technical signature. Nothing breaks, no threshold is crossed anywhere else on the panel, and the operational dashboards of a thesis failing this way are indistinguishable from those of a thesis succeeding. It is worth stating plainly that this is the mode in which the argument is most likely to be quietly wrong—not a fault in the machinery, but the discovery that a market which could have priced verification did not care to, because habit was cheaper, because a trusted intermediary was good enough, or because the buyer base the thesis assumed was never there. Section 13.3.10 argues that free verification deters misrepresentation before it is attempted; that argument requires the claimant to believe being caught would cost them something, and Condition A is the instrument that reads whether it does. This is the sense in which Red Line 14 tests the monetary mechanism and not merely the fee model: a wedge of zero means the market was offered symmetric knowledge and declined to pay for it.

Third-party checkability. The thesis's own principle applies with full force here: the party being scored must not set the scoring function (Section 17.8, Red Line 4). The reference panel, the workload-to-instance mapping, *and the verification cost share* γ must therefore all be maintained by someone other than the fee recipient, and an independent party must be able to re-derive ω from the published inputs and obtain the same number. γ deserves explicit mention because it is the easiest term to shade: it enters $\omega = s - \gamma$ negatively, so understating the cost of verification inflates the wedge and makes this red line harder to trip, and the fee recipient is the party with both the incentive and the telemetry. Its inputs — proving overhead and redundancy per unit, drawn from FER and VERIFYPRICE — must be published in the same

reproducible form as the panel, with the derivation runnable by an outside party against the raw receipts. A wedge computed by the entity collecting the wedge is a marketing figure, and that applies term by term.

Honest coverage caveat. P^{byp} is well defined only where a like-for-like unverified execution exists: matrix multiplication, inference, proof generation. It is *not* well defined for atomic settlement or media provenance, where there is no unverified version of the same product — the comparator there must be the all-in fee of a custodial or regulated rail delivering the same economic function. That is a rougher measurement with a weaker claim to like-for-like, and it must be reported separately rather than blended into a single index. An index that averages a clean comparison with a rough one inherits the rough one's error and hides it.

Effective prices, not list prices — and the direction of the residual bias. On-demand list pricing is not what hyperscaler and marketplace capacity actually transacts at. Committed-use discounts, reserved instances, and spot markets routinely clear well below list, so a panel scraping list prices measures a bypass channel nobody uses. That is why P^{byp} is specified above as the lowest *generally available* price across the published pricing modes rather than the on-demand quote, and why the panel spec must record which mode each reading came from.

Even so, a residual bias survives: the largest buyers transact at negotiated enterprise rates that are published nowhere, so the observable price remains an upper bound on the true alternative. The direction this pushes the red line is worth deriving rather than assuming, because it is not the direction one expects. Since $s = (P^{\text{prot}} - P^{\text{byp}})/P^{\text{byp}}$ falls as P^{byp} rises, an overstated P^{byp} *understates* both s and γ, and understates s by more, because P^{prot} exceeds the per-unit verification cost. Measured ω is therefore too small and Condition A trips *earlier* than the truth warrants. The reading is accordingly published as an **upper bound on breach**: a measured $\omega < \tau$ may overstate the case against the thesis and must be re-derived against negotiated pricing before the 90-day clock is treated as running, whereas a measured $\omega > \tau$ is a pass earned against a conservative comparator.

That is the benign direction, which is precisely why the term that runs the other way needs naming. Understating γ inflates ω and makes the red line harder to trip, and γ is the one input sourced from the fee recipient's own telemetry. The custody requirement above exists for that reason and is the load-bearing safeguard here, not the panel methodology.

Red Line 14 and its neighbours. Three distinctions worth stating precisely, because each is easy to collapse:

- **The price series is not VERIFYPRICE.** VERIFYPRICE is a *cost* SLO — the real-resource cost of checking a claim on reference hardware, deliberately exogenous to token price (Section 22.7) — and Red Line 1 reads on it. Red Line 14 does not: the quantity it reads on is a *price* series for delivered service, set by markets, which VERIFYPRICE is not and was never meant to be.

The two red lines therefore cannot be tripped by the same event. The qualification is that γ, one input to the wedge, does draw on VERIFYPRICE and FER telemetry for the cost of verification, which is the right source for a cost term; that is why γ carries the independent-reproducibility requirement above. Shared telemetry in one term, not a shared test.

- **This is not Red Line 6.** Red Line 6 asks whether fees follow usage — a direction. Red Line 14 asks whether there is a wedge worth charging at all — a magnitude (Section 13.3.4). A protocol can pass 6 while failing 14: native fees can track usage perfectly while the fee itself exceeds the differential value supplied.
- **This is not Red Line 5.** Fee coverage measures fees against the security budget. Red Line 14 measures the fee against what the market will bear. Coverage can be adequate on a base that is about to be competed away.

Red Line 15: Native Collateral–Capacity Spiral

Measured quantities. Publish native-asset peak-to-trough drawdown, collateral coverage against realized slashing exposure, operator exit and liquidation rates, stress-scenario DVC, collateral released into liquid float, and recovery of core service SLOs.

Condition. A material native-asset drawdown — defined here, and not to be redefined during an episode, as a peak-to-trough decline of $\geq 40\%$ in the native asset's price over a window of ≤ 90 days, measured in the workload unit-price index of Red Line 14 with the USD series published alongside — persists for at least 30 days below the trigger level; collateral coverage falls below realized slashing exposure; operator exit or forced liquidation exceeds its constitutional threshold; stress-adjusted DVC declines; released native collateral materially increases liquid float; and core service SLOs fail to recover within the 90-day remediation period. **The clauses are conjunctive by design** — each alone has an innocent explanation (a drawdown without coverage failure is a bear market; operator churn without DVC decline is turnover), and it is the joint deterioration that names the spiral. The thresholds must be set before the episode, reported for pure-native, mixed, and native-plus-insurance designs, and may not be relaxed while the clock is running.

Maturity gate. This line reads only after the Phase II entry gate of Section 29.3 is met; before that it is published as a watch indicator and trips nothing. A 40% drawdown in a Phase I asset with three operators is a venture repricing, and the collateral-coverage and DVC clauses have no series to read on.

Why it kills the thesis. The asset is being offered as insurance while its underwriting capacity disappears in the insured event. The loop

drawdown $\rightarrow$ more units required $\rightarrow$ operator exit $\rightarrow$ DVC decline $\rightarrow$ collateral release $\rightarrow$ liquid float inc

turns native value capture into a service-capacity contraction. A temporary drawdown is not the breach; the sustained joint deterioration and failed SLO recovery are.

Red Line 16: The Convenience-Yield Null

Maturity gate. This line does not read until the stack delivers measurable service at Phase III scale and a lending and derivatives complex exists around the base asset deep enough to quote continuously. Below that, the reading is published as *insufficient market* and trips nothing.

Measured quantities. The five-family convenience-yield vector of Section 26.8.6, published against the regime-pressure index R of Section 13.7.1: base-asset lending rates segmented by venue, term and borrower class; forward and perpetual basis net of the reference rate; and wrapper basis against the value of native holdings. Borrow demand is decomposed before the reading is taken — directional short interest, market-making inventory financing, and operational squeeze against locked float are not the premium, and an operational squeeze is a liquidity event reported as such. The *corridor band* is identified from the pre-committed pressure classifier and from independent service and access readings (Delivered Verified Capacity, VerifyReach, VERIFYSETTLE), never by locating where the yield series happens to peak.

Condition. Define the corridor increment

$$\theta_{16} \;=\; Y_{\text{band}} - Y_{\text{benign}},$$

where Y_{band} is the adjusted holder-service reading in the corridor band and Y_{benign} the adjusted reading in the benign comparison state. With the gate open, over a window of ≥ 8 consecutive quarters containing ≥ 2 observed regime transitions under the pre-committed state classifier (Section 29.2), the line trips when the pre-registered upper uncertainty bound on θ_{16} lies below the *minimum economically meaningful increment* — the smallest corridor increment the thesis would count as the mechanism operating, stated in basis points before the window opens. The band, the benign comparison state, the estimation specification, the confidence level, and the minimum increment are all pre-committed and published before the window opens. An interval that still admits the minimum increment is *inconclusive*, not success: it is published as such and the window extends.

Power, stated before the reading. A corridor increment estimated from a handful of quarterly observations will carry a wide interval whether or not the premium exists, and a red line that kills on low power is a red line that kills by construction. The published specification therefore carries a pre-registered power calculation: the minimum

increment, the sampling frequency actually available from the venues in the panel, and the resulting power at the window length. A bound below the minimum increment **kills only when the pre-registered power to detect that increment exceeds 80%.** Below that, the reading is published as *underpowered*, the window extends, and the line trips nothing — which is a statement about the instrument, not a reprieve for the claim. A window that never reaches 80% power is itself a published fact about how thin the market for this asset's regime-contingent service is.

Reading rules. A positive but flat reading — a holder-service level with no corridor increment — is not "priced at zero"; it means the claimed regime increment was not established, and it is published in those words. Extreme-pressure states in which delivery fails are not conditioned away: they stay in the full-stack report and on the availability red lines. The access experiment of Section 13.1 and these series answer different questions — whether holding changes what can be done, and whether a mature market prices the holder service — and neither substitutes for the other. Inventory financing and operational squeezes stay excluded; the preholding mechanism is not a route for relabelling operating demand as a premium.

Re-specified in v3.5, before the gate has opened. The condition previously read on a linear coefficient on R. The thesis's own Pressure–Capacity Corridor (Section 8.2) predicts an inverted U, and under symmetric sampling a linear fit to an inverted U can return a slope near zero while the mechanism is present. The band comparison tests the shape the thesis predicts. What now passes that could have tripped: an inverted U with a near-zero linear slope. What now trips that could have passed: a monotone series whose corridor increment is positive but below the minimum. What is unidentifiable: a window in which the classifier never places enough quarters in the band.

Why it kills the thesis: Section 13.3.9 shows that fee accrual is not a source of monetary premium, and Section 13.3.11 locates the premium to a holder-side service flow whose defining property is that it is *regime-contingent* — thin in benign states and largest in the corridor where substitutes weaken while the stack still works (why that flow should reach the base asset rather than Work Credits and Bitcoin is argued at Section 30.6.1, immediately below, and against conversion on the day at Section 13.1). A bound below the minimum increment is that property absent at the size the thesis needs: holders are charging no meaningful premium to part with units in exactly the states the asset is claimed to be for. What survives is a service asset with a cash-flow claim, a collateral lockup, and whatever pledgeability (Section 13.3.10) has been demonstrated — but the state-contingent premium the thesis argues for does not exist, and the argument for it must be withdrawn rather than deferred.

The null is the point. Section 26.8.6 already commits to publishing a null as a result.

This red line is that commitment given a clock, a threshold, and a consequence. An unsized quantity with no instrument is indistinguishable from an article of faith; an unsized quantity with an instrument and no trigger is an article of faith with a chart.

Insurance is priced when the storm is forecast, not after it lands. If the lending rate on this asset reads no higher in the corridor than in the calm, nobody was ever buying shelter — they were buying a security, and the thesis has spent six Parts describing a hedge its own holders decline to pay for. Either the rate moves with the weather or it does not, and an increment ruled out in the corridor is the market's verdict on the mechanism, delivered in basis points.

30.6.1 Why the Hedge Bid Reaches the Base Asset and Not Work Credits

The sharpest objection to the holder-side mechanism is internal to the thesis. A holder hedging loss of access to verified services already has an instrument for contingent access — the Work Credit — and a better exit hedge in Bitcoin. Rationally she holds both, and nothing above has said why the bid reaches the base asset at all.

The answer is that a Work Credit hedges the wrong state. It is typed and workload-specific; it is bounded by the Delivered Verified Capacity of the facilities behind it; it retires on redemption; and it is a claim on *specific* capacity. That capacity fails in exactly the states the hedge is for — the curtailment, embargo, or partition that removes a jurisdiction's provers removes the Work Credits written against them in the same moment. The instrument and the risk it would insure are correlated by construction. Bitcoin hedges the complementary state, exit, and hedges it well; it does not deliver a proof or a verified inference when the ordinary supplier will not. What remains is the one instrument that is bearer, workload-agnostic, does not retire, is the unit in which Work Credits are priced, and is what the protocol pays: the base asset. The hedge for "this service next quarter" is a Work Credit. The hedge for "access to *any* verified service in a repression state" is the base asset, and the holder-side mechanism attaches to the second alone.

That answers Work Credits, and it answers Bitcoin as an exit asset. It does not answer the third strategy: hold Bitcoin and buy the base asset on the day it is needed. Being the bearer, workload-agnostic unit gives a holder no claim on surviving capacity. The only advantage of holding it in advance is that a held balance stays spendable where acquiring one cannot be done in time, privately, or at all, while native settlement and the service path still work. Section 13.1 states that band, the access experiment that tests it, and the condition under which the mechanism is absent.

The corollary is conceded rather than argued around. If holders in fact prefer Bitcoin plus Work Credits, or Bitcoin plus conversion on the day — pricing exit and near-term service separately and leaving nothing for the agnostic claim — the access experiment finds no mate-

rial advantage, the convenience-yield series shows no corridor increment, and insiders never move obligations into the unit. That is precisely what the access experiment and Red Lines 16 and 18 detect, and the thesis treats it as its default null, not as an anomaly.

Red Line 17: Information-Sensitivity Reversion

What this line does not read on. Haircut *levels*. Cheap public verification removes one kind of information sensitivity — about unit quality: counterfeit, fraud, double-count, misreported delivery — and leaves the other kind untouched: sensitivity to the payoff, meaning future fee demand, regime, and governance (Section 13.3.10). The haircut level is set by the second kind, through realized volatility, and a 30–50% haircut on a volatile bearer asset is not a failure of anything argued here. What cheap verification predicts is narrower: that lenders who cannot disagree about what a unit *is* will disagree less about what to lend against it, and will re-converge faster after a shock, than lenders facing an opaque comparable of the same volatility. The prediction attaches in full force to *positions* with debt-like payoffs — over-collateralized, DVC-backed Work Credit positions; base-asset repo with volatility-set haircuts — and in attenuated form to the bare base asset.

Maturity gate. The line reads only when $N \geq 5$ independent lenders quote haircuts continuously on the base asset *and* on at least one opaque comparable over the whole window, at matched tenor and borrower class, with quoted depth published. Below that the reading is published as *insufficient panel* and trips nothing.

Measured quantities. Cross-lender haircut dispersion — the interquartile range of quoted haircuts across the pre-committed lender panel — on (i) the native base asset, (ii) pledged protocol positions, and (iii) an *opaque comparable set*: wrapped or custodial representations of the same asset, and tokenized commodity claims. Each series is residualized on trailing realized volatility and on quoted depth before comparison, under a specification pre-committed before the window opens.

Underwriting-model provenance. Each panel lender declares the provenance of its haircut model: independently maintained, a shared engine or vendor model, a common policy template, a common data dependency, or a mechanically copied schedule. Dispersion is published across the whole panel and across independently maintained approaches alone. Agreement generated by a common underwriting model is not counted as evidence that public verification removed an informational advantage; shared-model quotes stay published, with their dependence explicit in the inference. Five legally independent lenders running one decision system are one test, not five.

Condition. With the gate open, **(a)** the controlled dispersion on the native asset, across independently maintained approaches, is not lower than that of the opaque comparable set over ≥ 4 consecutive quarters; **or (b)** during a declared stress episode the native asset's dispersion widens *relative to* the comparables' — the ratio rises — or re-converges

more slowly, while receipt availability and VERIFYPRICE remain within SLO. Either reading, with no credible remediation path, measured against the pre-committed lender panel under the custody and reproducibility terms of Red Line 14. Levels are published alongside and are not the trip wire. A panel whose model dependence leaves fewer than five independently maintained approaches reads *unidentifiable*, not green, and a window that ends unidentifiable is itself published, as *underpowered* is for Red Line 16: it is a fact about how concentrated the lending market for this asset is.

Re-specified in v3.5, before the gate has opened. The relative design (native against opaque comparables, same lenders, matched tenor and borrower class) is unchanged. The provenance rule is added, and the *unidentifiable* reading with it. What now passes that could have tripped: a panel of five or more lenders whose quotes come from fewer than five independently maintained models, which now reads *unidentifiable*. What now trips that could have passed: a panel whose apparent agreement came from one shared model. What is unidentifiable: a panel without five independently maintained approaches.

Applied to Bitcoin. A formulation that read on dispersion declining as coverage rose would be tripped by Bitcoin — unit quality already perfectly verifiable, haircuts at 30–50% and dispersed across lenders — on the first reading. On this specification the test for Bitcoin compares native BTC against wrapped and exchange-custodied BTC quoted by the same lenders. It is a test Bitcoin may pass or fail; it is not one it fails by construction, and if it fails, the pledgeability mechanism is in doubt for this asset as well.

Why it kills the thesis: The asset is being priced on private information about specific units and positions, and public verification has not removed it. No-questions-asked acceptance has not formed, the pledgeability mechanism argued in Section 13.3.10 is absent, and cheap verification has bought operational convenience without buying moneyness. This red line reads on link 9 of Section 8.1, which the fee, capacity, and coverage lines do not.

Watch the disagreement, not the level. When independently underwriting lenders converge on one haircut, nobody believes there is anything left to find about the unit; when they scatter, each is pricing something it believes the others lack, and the scatter arrives before the level moves. Lenders running one model converge for a different reason, which is why the panel declares provenance. One reading of 2008 runs through the same margin: Gorton and Metrick ([Gorton & Metrick 2012]; Sources) document haircuts on structured collateral in the repo market rising from near zero to 40–45%, a repricing of what nobody could any longer check unit by unit. That is a level story, and this line does not read on levels; what it takes from the episode is *which* margin failed — unit quality, the one margin cheap verification addresses. A protocol whose receipts are public, whose verification is cheap, and whose lenders

disagree about what a position is worth no less than they disagree about its opaque twin has published everything and convinced nobody.

Reading note: borrower class. The pre-committed lender panel is segmented by borrower class—human, firm, agent policy (Section 26.8.9). The delegation model predicts that dispersion compresses first on agent borrowers, because an agent cannot do due diligence at speed and a lender to an agent cannot price what the agent will not disclose (Section 7.5). If dispersion compresses on agent panels and not on human ones, the mechanism is real and its first constituency is the predicted one. If it compresses on neither, the line reads null exactly as written. The segmentation changes no condition, gate, or threshold above.

Red Line 18: The Denomination Null

Maturity gate. Neither condition reads until free-choice contract volume in the relevant tier clears a published floor, stated here rather than deferred: trailing four-quarter free-choice notional in the tier, after the carve-out below, of at least 10% of settled protocol turnover, drawn from at least 50 distinct counterparties. Below that floor the reading is published as *insufficient base* and trips nothing; a red line that fires on an empty denominator is noise. Condition B carries a second floor of its own, given with it.

Accounting carve-out, applied before either condition is evaluated. Contracts whose quoting unit is fixed by an external reporting, tax, or regulatory requirement — functional-currency rules, auditor mandates, statutory reporting obligations — are excluded from the numerator and the denominator, and the excluded volume is published. Without this exclusion the line fires on accounting standards rather than on any judgment about the unit, and neither condition below is defensible without it.

Condition A — outsiders will not reckon in the unit. Tier II and Tier III free-choice denomination share, subsidy-adjusted, is flat or declining over four consecutive quarters while protocol physical throughput and native fee turnover grow over the same window; and the quote/settle matrix shows fiat-quoted, natively-settled contracts retaining or increasing their share of notional.

Why A weakens the thesis: the asset is a settlement rail rather than a unit of account. Counterparties transact *through* it and do not reckon *in* it — the configuration bitcoin has occupied for seventeen years while carrying monetary premium, which is why this is a downgrade of the monetary claim's completeness rather than of its existence. The thesis survives in reduced form: an accepted, pledgeable, state-contingently useful bearer asset that never became a coordination device, to be defended on Section 13.3.10 and Section 13.3.11 alone and restated as such.

Condition B — insiders will not reckon in the unit. B reads on *net* denomination, because an insider's revenue side is in the unit by construction and says nothing. Two

series, both after the carve-out: **(i)** the share of Tier I revenue quoted and settled in the unit, published as context; and **(ii)** the share of Tier I participants' freely negotiable *short-tenor* obligations denominated in the unit, computed over the *eligible* set: contracts whose counterparty would accept unit denomination on terms the insider could have taken, with the eligibility rule published. Tenor is the horizon over which the price term is fixed, not the date payment falls due: a five-year power contract invoiced monthly is long-tenor. Short-tenor means a price fixed for 90 days or less, a cohort definition fixed before measurement and not an empirical estimate. **Floor:** at Phase III entry (Section 29.3), $\geq 25\%$ of eligible short-tenor insider obligations unit-denominated. B trips, with the gate open, on failure to meet the floor — a network entering Phase III below it has tripped B on arrival — or on sustained failure to maintain it: (ii) below the floor for four consecutive quarters after entry. Insiders' sales of unit receipts within the settlement period against eligible obligations are published as supporting evidence of declined denomination opportunities, not as a separate trigger; ordinary treasury conversion is not condemned by itself. Long-tenor obligations — power purchase agreements, hardware financing, leases — are published as a separate series and carry no fatal threshold at this stage. The thesis claims only that they move last, if at all (Section 13.7.3), and makes no claim about when.

Why B kills the thesis: Tier I counterparties are protocol participants paid in the unit. Their costs are not: power, hardware, and payroll are fiat-denominated everywhere a proving fleet has yet been built, and the basis risk between the two is real. That basis is not a defect to be explained away; it is the choice B measures. An insider who carries the basis can resolve it in one of two directions — by moving obligations into the unit as counterparties permit, or by treating the unit as inventory to be sold on receipt. Tier I is the best-informed population that will ever exist about this asset, and no external accounting rule compels its choice once the carve-out is applied. If *that* population, offered the chance to denominate its short-tenor obligations in the unit it is paid in, keeps declining, the failure is not coordination among strangers. It is the people with the most information declining to reckon in the unit, which is the strongest available evidence that the unit is unfit to reckon in.

Why the two are held apart, and the standing counterexample. Bitcoin has a Tier I: miners, pools, hosting providers, and hardware vendors, paid in bitcoin and paying for power, ASIC financing, and payroll in fiat, with the industry's own benchmark — hashprice — quoted in dollars per petahash per day. On this specification Bitcoin would trip B on any reading now available: as far as public reporting shows, miners' short-tenor obligations are denominated in dollars, and selling coin on receipt to meet fiat costs is standard treasury practice, with a minority of treasury-holding miners as the excep-

tion. That is the standing counterexample, and it is stated at full strength rather than explained away: the strongest bearer asset in existence carries monetary premium while its own insiders decline to reckon in it. The severities still differ, because A describes a configuration Bitcoin has shown to be survivable — settling without denominating — and B describes one the thesis cannot survive, since a fee-denominated service economy whose own operators will not hold obligations in the unit has no constituency left to make the denomination argument. The thesis is therefore betting that a protocol which pays its insiders in the unit *and* prices their inputs in it can lift (ii) above the floor where Bitcoin's mining economy has not. If it cannot, B trips here for the same reason it would trip for Bitcoin, and the thesis grants itself no exemption. Either condition breaching trips the red line at its own severity; there is no requirement that they breach together.

Re-specified in v3.5, before the gate has opened. B(ii) previously read on long-dated obligations and tripped on a share that was merely non-increasing. Two defects followed: the line demanded early adoption in exactly the contracts the denomination theory says move last, and a stable share well above the floor tripped by remaining stable. B(ii) now reads on short-tenor obligations with a level-and-persistence rule, and sales on receipt are evidence rather than a trigger. What now passes that could have tripped: a stable insider share at or above the floor; insiders who convert receipts while denominating short-tenor obligations in the unit. What now trips that could have passed: a short-tenor share below the floor while long-dated obligations showed early movement. What is unidentifiable: an eligible short-tenor set below the volume gate. This is a substantive redesign of B, and it is published as one.

A unit people will spend but not quote in is a coupon. This test fails quietly, with every other board still green — throughput rising, fees arriving, corridors clearing, and every contract behind them written in dollars. That is bitcoin's seventeen-year position and it is survivable, which is why Condition A weakens the thesis rather than ending it. Condition B is the one to fear. It is not the world declining to reckon in the unit; it is the people who built it, who are paid in it, who post it as collateral, and who could write their own short-tenor obligations in it, choosing dollars anyway. Bitcoin's miners have made that choice for seventeen years. When the operators of a service economy priced in the unit make it too, there is nobody left to persuade.

Reading note: agent-issued contracts. Free-choice denomination is read on independently operated agents only. A unit an operator mandates for its own fleet is the mandated complement this line already excludes, and fleet-denominated volume is carved out before either condition is evaluated, with the carved-out volume published (Section 7.4). Delegation contracts denominated in the base asset by free choice may be the first place Tier II and Tier III share

becomes measurable; the falsifier is the fleet carve-out leaving nothing behind. The carve-out changes no condition, gate, or threshold above.

Red Line 16 and its neighbours. Three lines read on link 9 of Section 8.1, and they are not redundant — they test the three mechanisms the monetary claim actually rests on, and a network can fail any one while passing the others:

- **Red Line 16** tests the *holder-side flow*: does anyone pay more to hold in the corridor where substitutes weaken — and pay it for the base asset rather than for Work Credits, Bitcoin, or Bitcoin converted on the day (Section 30.6.1, Section 13.1)?
- **Red Line 17** tests *pledgeability*: will independently underwriting lenders disagree less about a unit whose quality anyone can check than about its opaque twin?
- **Red Line 18** tests *denomination*: will anyone reckon in it rather than merely settle in it?

The first fifteen red lines read principally on links 2 through 8 — the service path, the capacity, the fee accrual. Red Lines 9 and 12 arguably touch link 9 from the side, by reading on whether ownership and use take native form; none reads on the monetary mechanisms directly. Without the three lines above, that gap would be the sharpest available criticism of the falsification regime. They close it rather than argue it away.

A candidate line, published and not yet counted. The delegation economy (Chapter 7) raises one failure the eighteen lines do not read on: a native holder base whose addresses are many and whose operating rule is one. It is published here as a candidate so that its instrument can be built and its thresholds argued before it is asked to retire anything. It is not one of the eighteen, it trips nothing in this version, and every bracketed number in it is a placeholder.

Candidate Red Line 19: Principal Concentration

Status. Candidate. Not counted among the red lines; published as a watch reading only.
Maturity gate. Reads only after the Phase II entry gate of Section 29.3 is met, and only once the Principal Map classifies at least [50%] of each series it reads on as something other than *unknown*. A Phase I network's agent population is one operator by arithmetic, and a concentration ratio computed over a mostly unclassified base describes the classifier.
Condition. The Native Policy Concentration Ratio (Section 26.8.9) exceeds [50%] on any of prover collateral, corridor liquidity, or the self-custodied reserve cohort for ≥ 2 consecutive quarters; *or* the largest single operator class accounts for more than [70%] of agent-originated native fee turnover over the same period.
Instrument and clock. Principal Map (NPCR; agent-originated share by operator class). Two-quarter persistence, matching Red Lines 9, 10, 11, and 13.

Why it would weaken the thesis: Concentration of operating rules is a magnitude, as bypass is, and a step in one series can be innocent—a dominant open-source agent framework, freely adopted. But a constituency that is one template is a thinner anchor than link 8 of Section 8.1 asks for, and a thesis that counted it at face value would have to be restated with the holder anchor read principal-weighted rather than address-weighted.
Falsification of the line itself. If native positions are demonstrably diverse in policy—NPCR low—yet exit correlation during a declared stress episode is as high as if they were one template, policy concentration is the wrong proxy for behavioral concentration, and the line should be re-specified on exit correlation directly. Both series are published from the start so that the substitution can be made on evidence.

Red Line 3 reads on verification monoculture: one hardware profile, one vendor, one jurisdiction. Red Line 8 reads on the concentration of governance weight. Neither reads on the concentration of *operating rules* among native holders. A network can have diverse hardware, dispersed governance weight, and one rebalancing template, and the third is what decides a bad day.

Pattern: The red line is not "bad thing happens once." It is **"sustained breach + no recovery."** Isolated incidents with rapid remediation are expected in any complex system. Persistent degradation without response is thesis failure.

Condition	Response
Red line breached	Incident declared; governance must publish remediation plan within 14 days
Remediation fails after 90 days	Asset reclassified from "SoV candidate" to "speculative/experimental" on the public boards (Section 26.9); the reclassification is a published reading, not a statement by any body that speaks for the protocol, since Red Line 8 admits none
Multiple red lines breached simultaneously	Crisis mode; independent review commissioned; results published

Table 30.1: Red line response protocol. The clocks run on distinct triggers and do not nest: the 14-day publication clock starts at incident declaration; the 90-day remediation clock starts when the remediation plan is published (or at day 14 if none is); the two overlap only by construction of that sequencing, and each is reported with its own start date on the incident page. Neither clock resets on parameter change — the no-relaxation rule of Red Line 15 generalizes: thresholds and clocks are frozen for the duration of the episode they measure.

What happens at a red line: These are not punishments—they are **truth in advertising**. An asset that claims SoV properties must demonstrate them. If it cannot, it should stop claiming.

30.7 Market Realization Warnings

Red lines are protocol failures. The Market Realization Plane generates a second, weaker class of signal that must not be confused with them. A wrapper unwind can destroy price without damaging protocol function at all; treating that as thesis falsification would be as sloppy as treating a rally as confirmation. The distinction:

- **Monetary red lines** (above) mean the protocol's monetary claim has failed: verification breaks, settlement safety breaks, native value capture fails, issuance becomes discretionary, censorship routes dominate, or native use can be bypassed.
- **Market realization warnings** (below) mean *price has become an unreliable signal*. They do not necessarily kill the protocol; they suspend the right to cite price as evidence about it.

Warnings trigger when:

1. A statistically significant break appears in holder flow elasticity.
2. Net mechanical gain κ changes sign.
3. The Mechanical Pressure Ratio exceeds pre-declared market-depth bounds.
4. Wrapper creations dominate native spot demand.
5. Dealer swap capacity appears to bind.
6. Exposure migrates from swaps toward options.
7. Recursive leveraged wrappers appear.
8. Top wrappers or dealers exceed concentration thresholds.
9. The Wrapper–Native Growth Gap stays elevated.
10. Price rises while native fees and receipt volume fall.
11. Price falls while native monetary health improves.
12. Cross-jurisdiction hedges produce halt or settlement asymmetries.
13. Triad capacity is consumed at scale while the asset prices as a service and the monetary bid accrues elsewhere — **Service-Good Realization**, below.

The correct response to a warning is not intervention. It is **epistemic**: publish the warning, state which flow term is dominating, and stop using price as evidence in either direction until it clears. Governance must specifically not respond to warnings by inflating issuance, subsidizing price, or buying back supply — those are the reflexes of an entity managing a stock price, not operating a monetary constitution.

One warning is not merely epistemic. The twelfth item above is the last of the flow-mechanical warnings. The thirteenth is a different animal and is set out in full, because it is the observable form of the thesis losing the argument of Section 34.1.1 and because a warning that carries falsifying weight should not be left as a line in a list.

Warning 13: Service-Good Realization

Condition. All three clauses hold together, read across at least two distinct *regime-pressure episodes* — intervals in which the Stage A index R_t (Section 13.7.1) is elevated under criteria published before the reading rather than chosen after it:

1. **Capacity is being consumed at scale.** Receipt volume, verified units delivered, and private settlement volume grow through the episode; the capturable wedge stays above the realized take rate, so Red Line 14 does not fire; gross native fee turnover tracks usage, so Red Line 6 does not fire. *The protocol is working.*
2. **The insurance signature is absent.** Through the episode the asset exhibits nothing that distinguishes it from a claim on a service business: holding duration does not extend, the non-custodial share does not rise, the duration-neutral holding cohort of Section 29.3 does not grow, and demand moves with buyers' capacity budgets rather than with their exposure to the pressure. Read on the Value Capture and Agency Preservation boards (Section 26.9), not on a chart.
3. **The monetary bid accrues to Bitcoin.** Over the same episodes Bitcoin absorbs the flow the thesis predicts for a regime-contingent hedge, while the triad asset covaries with compute and IT spending rather than with regime pressure — and Work Credit inventory, not base-asset holding, is where any contingent-access demand shows up. That is the Bitcoin-plus-Work-Credits portfolio of Section 30.6.1, chosen by holders in the state the thesis said they would choose the base asset.

What it means. Not that the protocol failed — clause 1 says it did not. It means the regime-contingent convenience yield of Section 13.3.11 was not there, and that yield is the entire remaining basis for the monetary claim once Section 13.3.9 has shown that fees, burns, and collateral produce a cash-flow claim and a collateral lockup. The prescribed response is therefore *reclassification rather than retirement*: the asset is a verified-capacity service asset with a competitively priced fee stream and a collateral lockup. That is a real thing to be. It is not a store of value, and the protocol's communications must stop saying otherwise on the same 90-day clock the red-line protocol uses.

Why this remains a warning rather than another red line. Clause 3 reads on price, and reads it against another asset. Every condition in Section 30.6 is protocol-observable, and Corollary 13.1 forbids treating price as evidence about monetary adoption in either direction. A price-comparative red line would break both commitments. Red Line 15 instead covers the protocol-observable collateral–capacity spiral. What distinguishes this warning from its twelve neighbours is that clauses 1 and 2 carry the falsifying weight by themselves: capacity consumed at scale with no insurance signature, across repeated episodes, is the monetary claim failing whether or not anyone looks at Bitcoin. Clause 3 says where the premium went. It is corroboration, not the test.

Chapter 31

Implementation Sketches for Builders

Everything above is the blueprint. This chapter is about what to actually ship first; Chapter 32 says who ships it, on what horizon, and at what order of cost.

31.1 Minimum Viable Stack (MVS) Sequencing

Builders will ask: "What is the minimal sequence?" The answer matters because dependencies are real.

Priority	Component	What It Enables	Dependency
1	PIDL + verifier harness + anchoring	Telemetry cannot be faked	Foundation
2	One corridor kernel + VERIFYSETTLE	Non-custodial settlement	PIDL
3	One proof factory + VERIFYPRICE	Proof commodity market	PIDL
4	PAL SDK + PRK SDK	Developer surfaces	Corridor + factory
5	Neutral router + fairness tests	Open participation	SDKs + dashboards
6	PoUW pilots	Useful-work security	All above stable

Table 31.1: Minimum Viable Stack sequencing.

Why this order:

1. **Receipts first:** Without verifiable receipts, all claims are unpriced. Ship PIDL, reference verifiers, and dataset anchoring before anything else. This is the foundation that makes "no dashboards, no trust" real.
2. **One corridor before many:** A single refund-safe, non-custodial corridor with public VERIFYSETTLE is worth more than ten "coming soon" announcements. Ship BTC↔XMR first: its adaptor-signature protocol [Gugger 2020] is deployed and its refund semantics are understood (Section 23.4). BTC↔shielded-ZEC is gated on settlement work that does not yet exist and is a research item (Section 32.7), not a launch target. Refund safety means zero protocol-attributable losses with a bounded-time no-loss exit for every party, published with root-cause classification — not a 100% figure that trips on user liveness.
3. **One proof factory before PoUW:** Prove that canonical workloads can be produced, verified, and priced before adding consensus complexity. PROOF_2^{20} and MATMUL_4096_FP32_FREIVALDS

are good starting points.

4. **SDKs after primitives:** PAL and PRK are developer surfaces, not primitives. Ship them once the underlying corridors and proof factories are stable.
5. **Neutral routing after basic markets:** Fairness tests and admission metrics matter only once there is enough activity to measure. Do not over-engineer routing before you have provers to route.
6. **PoUW last:** Proof-of-useful-work consensus is the most complex piece. It requires stable proof factories, receipts, and measurement infrastructure. Do not attempt it until Priorities 1–5 are solid.

Hardware profile clarity: "Commodity GPUs with well-understood drivers" is an "open-ish hardware profile." To be explicit: this is **Profile Class: Partially Closed, High Familiarity**. GPUs are not verifiable machines in the same sense as open RTL—they are "known black boxes" whose behavior is well-characterized but not inspectable. Label profiles clearly so users understand the trust assumptions.

31.2 Layer-0 Verifiable Machines

Start with small, fully open "security evaluation" chips: key-storage/signing cores, simple TRNG/PUF arrays. For each tape-out, commit to a fixed sampling plan.

Minimal Layer-0/4 pattern:

- Pick one open-ish hardware profile (e.g., commodity GPUs).
- Define a hardware profile HID with chip family, driver versions, entropy tests, power metering.
- Stand up a proof factory cluster with reference prover binaries.
- Publish basic telemetry and VERIFYPRICE samples.

31.3 Privacy Rails: Making Settlement Safe and Boring

Standardize adaptor-signature swap flows. Take BTC↔XMR as the canonical deployed corridor and write down its message formats, time-lock conventions, and refund and punish procedures [COMIT 2021]. Treat BTC↔shielded-ZEC as the next corridor to *design*, not to standardize: there is no production protocol to write down yet.

Ship wallet UX with recovery, not bravado. Clear progress indicators, explicit timeouts, visible "abort and refund" buttons.

Document settlement latencies and success rates. Publish historical p50/p95 time-to-finality, success rates, and failure breakdowns.

31.4 Proof Factories: Receipts as a Service

The next implementation target is the Prove stage: turning raw work into portable receipts that any chain, app, or institution can consume.

Unify proof outputs under a receipt schema. Start by defining a compact, chain-agnostic receipt format (PIDL): the claim hash, workload ID, circuit or model hash, proof hash or transcript commitments, SLA tier, start/end timestamps, resource usage, and provider signatures. Whether the underlying artifact is a MatMul transcript, a logits fingerprint, or a SNARK, all of them fit into this envelope.

Expose a single programming surface. For developers, the SDK should offer one set of verbs: `declare_claim`, `request_proof`, `await_receipt`, `verify_receipt`. Underneath, it can target multiple proof systems and networks. The developer does not wire to "Platform X"; they attach their claim and policy to the SDK and let it find anyone who can satisfy the SLA.

Integrate multiple zkVMs and proof systems from day one. Monoculture is the enemy of both security and economics. The factory should be multi-ZK and multi-backend by construction: support at least one pairing-based SNARK, one STARK or FRI-based system, and one zkVM at launch, with a clear path to adding more.

Make telemetry a first-class product. Every proof request and receipt should emit metrics into the VERIFYPRICE Observatory: p50/p95 verifier times and costs for each workload, failure modes, queue depths per backend, variance under adversarial mixes.

31.5 Compute Consensus Pilots

With privacy rails and proof factories in place, you can begin to anchor consensus itself in useful work. Think in terms of reference pilots, not one chain to rule them all.

MatMul PoUW pilot (Duplex-style): Launch a devnet where the block-making puzzle is a MatMul instance drawn from a canonical distribution. Use the workload registry/WF-ABI to describe tasks; derive instances from header randomness; implement low-rank noise so that miners can return both a correct product and a succinct proof [Komargodski & Weinstein 2025]. Wire block validity to the presence of a valid MatMul receipt whose VERIFYPRICE parameters are below published thresholds. Instrument everything: $r(W)$ ratios, verifier times on consumer hardware, time-to-first-proof for new miners, and centralization metrics. State the devnet's purpose honestly: header-derived instances are unpredictable and therefore useless to any buyer, so this pilot measures consensus properties and verification cost, and does not by itself demonstrate useful work. The usefulness–unpredictability dilemma is the open problem of Section 22.4.1, and the devnet is the instrument for studying it, not evidence that it is solved.

Verified-inference pilot (Ambient-style PoL): In parallel, build a network where the work function is "serve model inferences with proofs of honesty." Commit to models and datasets; design logits-fingerprint schemes [Ambient 2024] that are simple, deterministic, and uniform; add randomized audits and peer-prediction-based slashing. Here full ZKML may be too expensive, so hybrid verification is acceptable, but only if it is exposed honestly: publish honest-output rates, audit coverage, and failure patterns as SLOs, not marketing.

zk-Proof-of-Work / receipt ledger pilot (Nockchain-style): Finally, either integrate with or prototype a chain whose primary job is to produce and timestamp zk proofs themselves [Nockchain 2025]. This network can serve as a common receipt ledger: proofs from many domains are anchored here with minimal metadata and availability guarantees.

These pilots need not be over-promised as final destinations. Their job in the early years is to act as laboratories: to establish that verification asymmetry holds in practice at scale, that decentralization can be preserved under useful-work mining, and that receipts remain cheap to check.

31.6 Developer Kit: Make "Import Proofs" the Default

The last implementation sketch is about developer experience. If every application needs a team of cryptographers to participate, nothing scales.

Claims & proving library: A library that plugs into common languages and frameworks and exposes a simple API: mark this function as "must be proven," mark this data stream as "must be provenance-tracked," specify acceptable backends and SLAs, and let PAL handle task generation and proof requests.

Settlement adapter: A component that connects to existing wallets and key-management setups, and that exposes "pay-for-proof" and "pay-for-compute" as simple intents: no manual swap logic, no bespoke bridges, just a predictable interface plus receipts, backed underneath by PRK and the atomic-payout kernel.

Verification module: Modules that can be compiled into smart contracts, browser bundles, and command-line tools, so that end-users and auditors can check receipts without standing up new infrastructure. `verify(receipt)` should be as embedded in the tooling as `log()` is today.

The work of the next few years is, in large part, to make these three pieces boring. When "import proofs" feels like "import TLS," and when "settle privately" feels like "call the payments API," the cypherpunk monetary stack stops being an argument and starts being the

default way serious systems are built.

Chapter 32

Kardashev Labs: Build and Measure

This is the one chapter in the thesis written for someone deciding whether to fund the work rather than whether to believe the argument. Chapter 31 says what to ship first. This chapter says who ships it, on what horizon, at what order of cost, and—the question an allocator asks before any of those—what exactly is being bought. It sits here rather than in the front matter because its claims depend on everything before it.

32.1 What Is Being Bought: Stack First, Coin Optional

The thesis has conceded, in its own voice, four things an allocator will have noticed. Bitcoin's buyerless work function is the monetarily superior design and nothing here improves on it (Section 34.1.1). The closest utility-token precedents reached low-single-digit fee coverage and never acquired monetary behaviour (Section 34.1.3). The fee base on which Phase III depends is expected to become measurable only on the far side of a financial clearing event the thesis does not predict and cannot time (Section 29.3). And the residual monetary premium—the three mechanisms of Section 13.3.10, Section 13.3.11, and Section 13.7.3—is unsized, without precedent, and instrumented by red lines that the thesis treats as its default null (Section 30.6.1).

Read together, those concessions fix what is fundable as infrastructure and what remains an experiment.

The Investment Thesis, Stated Once

Fundable as public infrastructure: the verifiable stack as public infrastructure with a measurement institution attached—the VERIFYPRICE Observatory and its adversarial harness, VerifyReach, VERIFYSETTLE on one deployed non-custodial corridor, the PIDL receipt schema and SDKs, the Facility Capacity Receipt schema, and VerifyFlow wrapper telemetry. Every one of these has a prospective buyer that is not a token holder: auditors, regulators, procurement desks, allocators pricing existing wrappers, and operators who need their capacity claims believed. Several could be sold before any base asset exists; whether any is sold is the "build without buyers" test below.

An option to learn, not an underwriting: whether a separate bearer base asset behind that stack earns monetary premium. Funders hold no claim on that asset. The thesis rates this conditional, gives it an empty base rate for fee-stream claims, and publishes eighteen conditions under which it is retired. Nothing in the lab's budget depends on it being true, and nothing in the lab's deliverables is worthless if it is false.
What the lab is therefore not: a token launch, a protocol foundation, or a proxy for buying the base asset early. Kardashev Labs issues no monetary instrument and holds no protocol treasury. If a base asset is ever issued, it is issued by a protocol whose constitution (Chapter 25) the lab may have drafted and must be able to fail.

That statement costs the thesis nothing it has not already conceded, and it is the only form in which a sophisticated allocator would fund the work. The monetary experiment rides on top of infrastructure that pays for itself in measurement, or it does not ride at all.

32.2 What Kardashev Labs Is

The name is a statement of scope, not of ambition: a civilisation's position on the Kardashev scale is a measurement of the energy it can *verifiably* command, and this thesis has spent six Parts arguing that verification is the binding constraint on the next monetary object. The lab's model is Bell Labs in one respect only—research and engineering under one roof, with the engineering disciplined by a customer—and it departs from Bell Labs in the respect that matters most here: its outputs are public goods by default, and its measurement outputs are published under the reproducibility contract of Section 26.6 whether or not they flatter the thesis.

Three functions, in priority order.

1. **Measure.** Run the observatory: VERIFYPRICE, VerifyReach, VERIFYSETTLE, VerifyFlow, and the Native Monetary Buyer Map, with signed multi-operator datasets anchored publicly. This is the lab's first product and its permanent one. It is also the function that can be sold on day one: VerifyFlow against existing bitcoin wrappers requires no protocol to exist.
2. **Build.** Ship the Minimum Viable Stack of Section 31.1 as reference implementations under open licences: PIDL and the verifier harness, one BTC↔XMR corridor kernel, one multi-backend proof factory, the PAL and PRK SDKs, and the neutral-router fairness tests.
3. **Research.** Attack the open problems the thesis has been honest enough to name as open (Section 32.7), and publish negative results with the same prominence as positive ones.

Relationship to any protocol. The lab writes specifications, reference implementations, and measurement; it does not operate the protocol, hold its treasury, or receive its fees. Any protocol adopting the stack pays the lab, if at all, as one customer among several for measurement

and audit services, under published terms, so that the party being scored does not fund the scorer in a way it can withdraw (Section 17.8 applies the same rule to $\mathcal{O}_s$ weights). The lab's independence from the protocol is what makes its dashboards worth reading.

32.3 Deliverables on a Dated Horizon

Chapter 29 omits calendar horizons for *adoption*, and it is right to: adoption is gated by metrics, not by quarters. A build programme is different. An allocator cannot underwrite "when the gates clear." The table below therefore commits to dated deliverables from a funding start T_0, counted in quarters (T_0+n is n quarters after funding), and it is explicit about which phase of Section 29.3 each deliverable belongs to. Every one of them is a Phase I artifact. That is not a hedge; it is the consequence of Section 29.3: Phase III is not reachable on a 36-month horizon under the thesis's own model, and a plan that claimed otherwise would contradict the document it is attached to.

Table 32.1: Kardashev Labs deliverables from funding start T_0, by quarter. Every row is a Phase I artifact. Acceptance tests are public and adversarial by construction; a row that fails its test is reported as failed.

Quarter	Deliverable	Acceptance test (public)	MVS step
T_0+1	PIDL v1 schema; reference verifier harness on three canonical workloads (PROOF_2^{20}, MATMUL_4096_FP32_FREIVALDS [Freivalds 1977], one settlement workload) plus a verifier stub for the SNARK SKU, which has no prover to measure; adversarial acceptance suite	Harness reproduces published VERIFYPRICE on two independent operators' hardware within the reproducibility contract, for every SKU that has a prover	1
T_0+2	VERIFYPRICE Observatory live with signed, anchored datasets; Laptop-Class and Server-Class reference verifiers published as bills of materials	Six weeks of continuous public series; first adversarial-mix report	1, 3
T_0+3	VerifyFlow v1 against existing bitcoin wrappers: κ, WRR, Q^{RC}, MPR published weekly	Reproducible from public ETF data by a third party; first paying subscriber	—
T_0+4	BTC↔XMR corridor kernel with refund-first wallet UX; VERIFYSETTLE live	$\geq$ 1,000 mainnet swaps; protocol-attributable refund failures $= 0$; overall success $\geq 95\%$; p95 time-to-finality published	2

Quarter	Deliverable	Acceptance test (public)	MVS step
T_0+5	Multi-backend proof factory (one pairing SNARK, one STARK, one zkVM) issuing PIDL receipts; VERIFYPRICE per backend	Three months of receipts; $r(W)$ inside modality band for every registered workload	3
T_0+6	PAL and PRK SDKs; `verify(receipt)` modules for contract, browser, CLI; first reference application (procurement or media provenance) in production with a named counterparty	External team ships against the SDK without lab assistance; receipts consumed by an auditor	4
T_0+6	Access experiment (Section 13.1) pre-registered and run as an analog on an existing network whose fees must be paid in its own asset, against Bitcoin-funded procurement including provider-side conversion	Protocol and equal-budget strategies published before the run; results published whichever strategy wins	—
T_0+7	VerifyReach v1 on OONI methodology [OONI 2012–] across ≥ 50 ASNs; transport fallbacks (look-like-TLS, Snowflake, refraction [Frolov et al. 2019]) measured	Six months of public reachability series; first censored-region report	—
T_0+8	Neutral router with published fairness tests; FCR schema v1 with two facilities reporting	Fairness tests pass on a mixed prover set; FCR fields independently verified for one facility	5
T_0+9 to +12	Layer-0 sampling programme pilot on one L0-A profile (counterfeit, binning, firmware-substitution targets); first open L0-C components (RoT, RNG, metering) taped out on an available open PDK; MatMul-PoUW devnet with the usefulness dilemma instrumented rather than assumed away	Published sampling report with stated sensitivity; devnet centralization and $r(W)$ series	6

Two rows deserve comment. A BTC↔ZEC corridor, the natural co-equal first target, is absent, because shielded-ZEC adaptor-signature swaps do not exist and transparent-ZEC HTLCs forfeit the privacy the corridor is for (Section 23.4); it returns when that work is done, and the lab's research agenda includes it. And the PoUW devnet is last and is labelled a laboratory,

because Section 32.7 names the problem it exists to study.

32.4 Team Shape

The lab is small enough that its director knows every result personally, and it is built in this order:

- **Measurement (first hires, 4–5):** two measurement engineers who own the observatory and the reproducibility contract; one data engineer for anchoring, signing, and publication; one quantitative analyst for VerifyFlow and the Buyer Map; one adversarial tester whose job is to break the harness.
- **Proof systems and workloads (3–4):** engineers with production experience in at least two proof systems; one with numerical-analysis depth for acceptance bounds and workload registration.
- **Settlement and privacy (2–3):** corridor and wallet engineers with adaptor-signature and shielded-protocol experience; one security reviewer.
- **Network and distribution (2):** transport and censorship-measurement engineering; software-update and transparency-log engineering.
- **Hardware (1–2, from T_0+6):** an open-silicon engineer for the L0-C components and a sampling-programme lead with a hardware-security background.
- **Direction (2):** a director who has run a measurement or standards organisation, and a research lead who has published negative results.

Governance, legal, and policy counsel are retained, not hired, until Chapter 27 says otherwise.

32.5 Planning Envelope

This section states the *structure* of the envelope and deliberately no figures. A CC BY document is the wrong place for a budget: numbers in it anchor every later negotiation, age within a year, and turn a thesis into a pitch. A costed plan—roles at loaded cost by month, capital items by class, vendor quotes where they exist, and a sensitivity table—is held as a separate document for qualified counterparties, and its existence rather than its contents is what this chapter records.

The envelope has six lines, and each is driven by something the thesis already specifies:

- **People.** The team of Section 32.4, at senior engineering and research compensation, fully loaded. The dominant line by a wide margin; its driver is the hiring order, not the headcount.
- **Compute and proving capital.** A GPU cluster for the proof factory and canonical-workload benchmarking, and reference-verifier fleets in the Laptop and Server classes of Section 22.6.

Driven by the workload registry and by the proving-cost series of Appendix E.6, which is why that series is published before the second tranche.

- **Measurement operations.** Multi-operator observatory hosting, anchoring, VerifyReach vantage points across at least fifty ASNs, and data publication. Driven by the operator count the reproducibility contract requires.
- **Security review.** External review of the corridor kernel, SDKs, and harness; bug bounties. Driven by the number of shipped surfaces.
- **Open silicon** (from T_0+6). Shuttle runs on an available open PDK for the RoT, RNG, and metering components of Appendix E.6, and lab time for the sampling pilot on one profile. Driven by shuttle availability, which the lab does not control.
- **Legal, policy, and administration.** Entity formation, licensing, and counsel retained under Chapter 27.

Excluded by design: any hardware-sampling programme at the scale Section 17.4 contemplates for a live network, which is funded from protocol fees and not by the lab; any protocol treasury; and any base-asset purchase.

Tranches and gates. The plan is funded in tranches gated by the deliverables of Table 32.1 rather than by the calendar, which is the thesis's own rule applied to the institution proposing to test it. The first tranche runs to T_0+6 and is released against the receipt schema, the harness, and the first VerifyPrice publication; the second runs to T_0+12 against the corridor, the observatory, and the first proving-cost reading; the extension to 36 months is not offered unless the "build without buyers" condition of Section 32.8 has been passed. An allocator who wants a single cheque for three years is being asked to fund a forecast, and this document does not publish forecasts.

Revenue against the envelope. The lab is not designed to be self-funding on this horizon, and it should not pretend to be. Three revenue lines are real and should be pursued from T_0+3: VerifyFlow subscriptions from allocators pricing existing wrappers; measurement and audit engagements for operators who need FCR and VERIFYPRICE attestations; and reference-application integration work with named counterparties. Together they are unlikely to cover more than a fifth to a third of the 36-month envelope, and the remainder is grant, consortium, or strategic capital that is buying public goods and an option, in that order.

32.6 Entity, Licensing, and What the Lab Owns

The lab is best housed in a non-profit or public-benefit entity whose charter forbids it from issuing a monetary instrument, holding a protocol treasury, or taking fees that any single protocol can withdraw. Reference implementations ship under permissive open-source licences;

measurement datasets ship under an open-data licence with the signing and anchoring rules of Section 26.5; specifications are public. What the lab owns is its reputation for publishing numbers that do not flatter it, and that asset is destroyed by the first exception.

Funders receive what a funder of a measurement institution receives: governance seats on the entity, priority access to data and engineering time under published terms, and the option to fund any protocol the lab's work makes possible on the same footing as anyone else. They do not receive tokens, because there are none to give.

32.7 Research Agenda: The Problems Named as Open

A build lab that only builds is a contractor. The following problems are the ones this thesis has, in the course of correcting itself, identified as unsolved, and they are the reason the institution needs a research function. Each is stated with the chapter that names it and the result that would count as progress.

1. **Usefulness versus unpredictability in proof-of-useful-work** (Section 22.4.1). Header-seeded inputs are useless to buyers; client-supplied inputs are precomputable. Progress: a hybrid design with a stated, tested collusion bound between client and miner.
2. **Verifiable accelerator compute** (Section 17.4). Generic GPU work cannot yet be proof-wrapped at acceptable overhead. Progress: the proving-cost series of Appendix E.6—cost per 10^9 cycles or constraints on the reference prover, measured quarterly, with the proving overhead per SKU beside it—with the modality bet stated as a slope and its eight-quarter kill condition, and the class of workloads for which M1 verification is economic at each point on the curve.
3. **Detection sensitivity of hardware sampling** (Section 17.4). The sampling bound assumes a sensitivity nobody has measured for leading-node silicon. Progress: a published σ for each threat class, beginning with the ones sampling can actually catch.
4. **Shielded cross-chain settlement** (Section 23.4). BTC↔shielded-ZEC atomic swaps with bounded-time no-loss exit do not exist. Progress: a construction, or a proof that one requires protocol changes on the ZEC side.
5. **Delivered Verified Capacity as a computable object** (Section 17.9). The flow model is not yet unit-coherent across the energy-to-proof chain. Progress: a generalized-flow formulation with a published graph-construction procedure, run against real FCR data.
6. **Whether cheap verification moves collateral terms at all** (Section 13.3.10). The scoped prediction—lower haircut dispersion and faster convergence, not lower levels—has never been tested on any asset. Progress: the study on bitcoin's own record, native units against wrapped, run before any base asset exists, so that Red Line 17 has a baseline.
7. **Whether holding in advance beats converting on the day** (Section 13.1). The holder service rests on a balance staying spendable when acquisition fails, and the strongest alternative is

Bitcoin plus provider-side conversion. Progress: the analog run of Table 32.1, then the same protocol on the reference network once one exists.

8. **The base-rate question** (Chapter 35). No equity-like claim on a fee stream is known to have become money. Progress: a documented negative, or the first counterexample, from the historical record.

Items 6 and 7 are the ones the lab should run first: both cost little, neither needs a protocol of the lab's own, and each could retire part of the monetary claim — pledgeability and the preholding service respectively — before a dollar is spent building for it.

32.8 How the Lab Fails

An institution built to publish kill conditions for a thesis should publish its own.

- **Measurement capture.** If the observatory's datasets cannot be reproduced by an operator the lab does not control for two consecutive quarters, the lab has become the thing Red Line 4 describes and should be wound down or re-chartered.
- **Build without buyers.** If, by T_0+8, no counterparty outside the lab consumes a receipt, subscribes to VerifyFlow, or ships against the SDK, the infrastructure has no demonstrated demand and the envelope should not be extended to 36 months.
- **Research without results.** If none of the eight problems above has a published positive or negative result by T_0+12, the research function is not earning its share of the envelope.
- **Capture by a protocol.** If any single protocol's payments exceed a published share of the lab's revenue, or if the lab's staff hold a base asset the lab measures without disclosure, the independence that makes the dashboards worth reading is gone.

None of these is a red line for the thesis. They are red lines for the institution proposed to test it, and an allocator should hold the lab to them with the same lack of sentiment the lab is asked to bring to the thesis.

Chapter 33

The Closed Sovereign Stack

Everything argued so far has been argued in the abstract: when soft guarantees weaken, value migrates toward hardened, verifiable, vertically coherent capacity, and five parts describe how to build such a stack.

A reasonable reader will ask whether anyone is actually doing this, and whether it works.

The answer is yes, at civilizational scale, by a state, with the politics inverted. That case deserves a chapter of its own—not as a prediction, an endorsement, or a geopolitical forecast, but because it is the closest thing this thesis has to an existence proof and to a counterexample simultaneously.

Epistemic Status of This Chapter

The material here is of three kinds and they are not interchangeable.

Primary sources, used for what they state: published policy speeches and central-bank staff research. Where this chapter quotes an official, the quotation is verbatim and the claim is about stated doctrine, not about outcomes.

Analyst commentary, principally Doomberg's *Fire Horse* [Doomberg 2026a] presentation and macro work by Luke Gromen [Gromen 2025–26]. Neither is peer-reviewed. Figures attributed to them are *their* accounting rather than independent measurement, and are used to characterize a strategy rather than to establish a quantity.

Scenarios, which are quarantined in Appendix I together with the claims from the same sources that did not survive checking.

Nothing in Section 30.6 depends on this chapter, and no red line, threshold, or measurement contract is derived from it. What the chapter supports is corroborative rather than foundational: Premise 1 is argued on its own grounds in Chapter 2 and Chapter 3, and this chapter corroborates it by observation. If the figures are wrong in detail, the argument here weakens as illustration and the rest of the document is unaffected. Nothing in Section 30.6 depends on it.

33.1 A State That Already Assumes Soft Guarantees Have Failed

Chapter 5 builds an adversary model from the premise that institutional promises are becoming conditional. That premise is usually defended by argument. It can also be defended by observation: some actors are already behaving as though it were true, and paying enormous costs to do so.

China's energy and industrial strategy is best read as a sustained bet against the reliability of external guarantees [Doomberg 2026a]. The recurring pattern is not accumulation of any single resource but maximization of the number of independent pathways from a controlled input to a required output:

- Domestic coal retained at scale, not because it is clean or cheap in the long run, but because it can be mined inside the system and therefore cannot be interdicted at sea.
- Refining capacity built beyond immediate domestic need, with feedstock flexibility across crude grades, LPG, and ethane.
- Coal-to-liquids, coal-to-gas, and coal-to-chemicals facilities that are difficult to justify on conventional return-on-capital grounds.
- Nuclear technology licensed from suppliers in rival blocs — Westinghouse, Framatome, Rosatom, AECL — plus indigenous designs, so that no single provider controls the program.
- Caution about additional pipeline dependence even on friendly suppliers.
- Statutory minimum coal stockpiles—which purchase not fuel but *time*, the interval during which the system can adapt.
- Electrification of transport, converting an imported-oil vulnerability into demand for electricity and batteries produced domestically.

The organizing principle is conversion optionality: many paths, deliberately redundant, accepting poor measured returns in exchange for freedom of action. Several of these investments are straightforwardly value-destroying under normal market conditions. That is the point. Their value is state-contingent, concentrated in exactly the scenarios where ordinary market relationships stop clearing.

> A synthetic-fuel plant that loses money nineteen years in twenty and prevents a transportation crisis in the twentieth is a bad investment and an excellent insurance policy. Discounted cash flow prices the first sentence and not the second.

Resilience buys time, and time is a weapon. The stockpile item above is easy to read as a defensive measure, and reading it that way understates it. A statutory minimum inventory purchases an interval. What the holder does with that interval is a separate question, and the answer is not necessarily defensive.

A system that can defer, substitute, ration, or absorb a disruption acquires something more useful than protection: it acquires influence over *when the disruption becomes binding*. It can decline to bid during a scramble and replenish afterward, at lower prices. It can wait out an adversary's electoral cycle, funding round, or refinancing calendar. It can allow a shock to run knowing that the counterparty with the shorter inventory and the more leverage-sensitive balance sheet will have to move first. Redundancy converts into tempo, and tempo converts into bargaining power:

$$\text{resilience} \rightarrow \text{time} \rightarrow \text{strategic initiative}$$

Strategic Tempo

The conversion of physical resilience—stockpiles, spare capacity, fuel-switching and conversion flexibility—into control over *when* a disruption becomes binding, and thence into bargaining power. Tempo is the return on redundancy that conventional cost accounting does not price.

This is a strictly stronger claim than the option-value argument of Section 17.10, and it compounds with it. Option value says the redundant facility pays off in the bad state. Tempo says its existence changes which bad states occur, and when—because a counterparty that knows you can wait negotiates differently than one who knows you cannot. For Layer 0 this means the sovereign optionality index of Section 17.8 is measuring something whose value is understated by its own construction: the index counts pathways, and pathways buy time, and time is worth more than the pathways cost.

This is the same structure as the option value argued in Section 17.10: facility value contains commercial, contingency, deterrence, and bargaining components, and conventional analysis captures only the first. A country that can survive an embargo has more leverage *before* the embargo, whether or not the capacity is ever used.

Why this matters for Premise 1. The thesis argues in Section 1.2 that soft guarantees are weakening. A state spending decades and enormous sums to become less dependent on external promises is expressing the same judgment through capital allocation rather than through argument. Revealed preference at sovereign scale corroborates that premise without establishing it: the argument for Premise 1 is made in Chapter 2 and Chapter 3 and does not rest here. What this chapter adds is that the judgment is not only this thesis's, and it is available without requiring any normative agreement with the actor.

33.2 Trust Minimization Applied to Matter

The cypherpunk instruction is familiar in its slogan form, and this thesis corrects it in Chapter 1: the working instruction is *verify in order to trust*. Stated that way:

> Verify in order to trust.

The strategic instruction described above is structurally identical, once "verify" is read as "retain an independent means of knowing":

> Don't trust; possess, duplicate, stockpile, or retain an alternative conversion pathway.

Both are responses to the same problem—a promise is weaker than a possession—and both accept significant efficiency losses to reduce dependence on counterparty good behavior:

- A contractual assurance that oil will arrive is weaker than an inventory of oil.
- An assurance that a vendor will keep supplying technology is weaker than domestic production.
- An assurance that an ally stays friendly is weaker than diversified supply.
- An assurance that trade lanes remain open is weaker than domestic conversion capacity.

Every line above is a Layer 0 argument (Chapter 17) written in molecules instead of silicon. The recognition is useful in both directions: it shows that the thesis's physical premises are not eccentric, and that trust minimization is politically neutral machinery that serves whoever builds it.

One precision on both slogans. Each instruction above is stated in its familiar compressed form, and each, read literally, prescribes a permanent distrust that neither an economy nor a state could afford to practice. Neither is meant that way. A stockpile is not held so that trade never happens; it is held so that trade can be entered into without the counterparty's continued good behavior being load-bearing. Cheap verification does the same work by the same route (Section 13.3.10): its purpose is not to abolish delegation but to make delegation survivable, by ensuring that the party who would exploit it knows the exploitation is checkable. **What both instructions minimize is not trust. It is unpriced dependence.** The distinction matters for the comparison this chapter is drawing, because it is the point on which the two stacks diverge rather than converge: both reduce unpriced dependence for the entity that builds them, and only one of them reduces it for the participant.

33.3 The Mirror: The Same Doctrine, Arriving in the West

The preceding sections treat the closed stack as a foreign object—coherent, instructive, and elsewhere. That framing has an expiry date, and it has probably passed.

On 23 June 2026, the U.S. Secretary of the Treasury told the Economic Club of New York that "economic security begins with national capacity," identified semiconductors, artificial intelligence, quantum computing, advanced manufacturing, shipbuilding, critical minerals, and pharmaceuticals as the industries that would define the next century, and argued that supply chains must be able to withstand coercion and crisis rather than merely minimize cost. The speech quotes Hamilton directly on enlarging the sphere of domestic commerce. The Hamiltonian framing is not an outside interpretation imposed on the policy; it is the policy's account of itself. (Source: U.S. Department of the Treasury, press release sb0539, "American Economic Statecraft in the 21st Century," 23 June 2026 [Bessent 2026]; Sources. Quotations here are from the published text.)

Read against Section 33.1, this is recognizably the same doctrine: capacity valued above cost, resilience valued above efficiency, strategic continuity valued above return on capital.

What the speech does not say, and the discipline of noting it. It does not advocate gold settlement. It does not propose capital controls. Its stated posture is "open to the world while anchored at home," and it treats digital assets, stablecoins, tokenization, and payment standards as instruments for extending dollar financial leadership rather than for retreating from it. Commentary that reads this material as a program of monetary retreat is reading in something that is not present, and Appendix I.4 lists the specific claims of that kind that this thesis declines to use. The accurate description is selective protection plus industrial subsidy plus investment screening plus dollar-denominated digital rails—which is a profound regime change without being the one the more excitable commentary describes.

The asymmetry. What makes the Western version harder than the original is not doctrine but the balance sheet it must be financed from.

The closed stack of Section 33.1 was built while suppressing household consumption, directing credit administratively, and tolerating poor measured returns for decades. The Western attempt begins from the opposite configuration and must preserve, simultaneously: high asset valuations, high corporate margins, entitlement commitments, low consumer prices, an open capital account, a reserve currency, and low inflation. These cannot all be maximized at once, and Section 3.2 explains why the constraint binds harder than it looks—the asset valuations are not merely a political preference but a fiscal input.

> The West is attempting to move from financial sovereignty back toward physical sovereignty, financed through a balance sheet optimized for the regime it is leaving.

There is a second, independent reason the Western attempt is hard, and it comes from the buyer side rather than the fiscal side (Section 3.1.3). China's directed banking system is a **duration warehouse**: it can be compelled to hold the interval between pouring concrete and producing power even when expected private returns are poor. That creates misallocation and hidden losses. It also enables projects whose strategic option value exceeds their commercial value. The United States has relied on foreign insurers, liability-driven pensions, retirement defaults, market-value indexes, and leveraged relative-value intermediaries. Section 3.1.3 argues that this buyer architecture is not a substitute for the old liability-matched bid. The physical reconstruction therefore arrives at the same moment the financial system is becoming less willing to own maturity. The closed stack solves time by command. The open financial system is currently solving it by shortening the public book into bills—which transfers duration onto the rollover calendar rather than extinguishing it (Section 3.2).

The reshoring trilemma. Reindustrialization can be fast, cheap, or excellent; a system may choose two [Gromen 2025–26]. Strategic competition demands fast. Bond yields and inflation demand cheap. The security rationale is void unless the result is excellent. After decades of deindustrialization the binding constraints are not primarily financial—skilled trades, engineering depth, machine tools, grid interconnection, transformer and switchgear supply, permitting throughput, supply-chain density, and industrial project management are all lead-time constrained, and capital cannot compress lead times past a point. The realistic expectation is therefore reshoring that happens, and takes longer and costs more than its advocates project.

Reindustrialization is a restructuring of claims. This is the part with monetary consequences, and it is the reason the chapter belongs in this document rather than in a policy annex.

Converting a financialized economy into a productive one means turning financial claims into factories, generation, transmission, mines, machine tools, trained labor, and fabrication capacity. If an economy has issued more long-duration claims on future output than its physical system can honor *while also* funding that conversion, the transition must impair some claims. The impairment can arrive as inflation, negative real rates, taxation, margin compression, capital controls, or explicit restructuring, and in practice it arrives as a combination.

> The industrial transition is also a liability-side restructuring. Someone's claim on future output is not going to be honored at present value, and which someone is a political question with no automatic answer.

The tempting simplification is generational—bondholding retirees versus wage-earning

young. It does not survive contact with the distribution. Inflation harms renters, cash-poor workers, households without bargaining power, and anyone trying to acquire assets before wages adjust; older households hold equities, real estate, and operating businesses, not only bonds. The defensible statement is narrower:

> Reindustrialization redistributes away from the beneficiaries of asset-price inflation, cheap imports, suppressed labor shares, and long-duration financial claims—toward some combination of labor, domestic producers, strategic industries, and the state. Which combination is decided by policy design, not by the transition itself.

That last clause is the whole of the thesis's interest in the question. An industrial policy without broad ownership and wage gains does not dissolve a financial oligarchy; it substitutes a subsidized industrial one, and the participant is no better off for the change. This is Section 3.6 restated at national scale, and it is the reason Section 5.2 treats protective industrial policy as an adversary class rather than as a neutral backdrop.

33.4 Open Weights Are Not Open Sovereignty

The doctrine of Section 33.1 has extended to intelligence, and the extension is easy for a reader of this thesis to misread as good news.

Openly released model weights are now within reach of the proprietary frontier. The Kimi K3 release [Kimi Team 2026], a 2.8-trillion-parameter mixture-of-experts model published with full weights, is presented by its authors as an open frontier system; the same paper reports that it trails the strongest proprietary models overall while leading the other systems in its evaluation suite. That is the claim this thesis makes and no more: open-weight capability is close enough to the frontier to be economically consequential.

The consequence is a familiar strategic pattern applied to a new good. In solar, batteries, electric vehicles, steel, and telecommunications equipment, the objective was never to capture the highest margin. It was to make adequate capability abundant at a price beneath the level competitors' capital structures require. Software weights are unusually well suited to that strategy, because once trained they can be distributed at approximately zero marginal cost. A sector priced for durable scarcity rents does not need to be defeated technically; it needs only to have the scarcity assumption questioned credibly. Whether this is designed as economic strategy is not observable and is not claimed—Appendix I.4 sets out why intent claims here are unusable. The effect does not depend on the motive.

The category error. None of this makes open-weight models an ally of the open stack, and the temptation to treat them as one should be resisted firmly.

Open Technology vs. Open Sovereignty

Open-source or open-weight software is sovereign only when its users can independently obtain the power, hardware, data, communications, privacy, and settlement required to operate it. Openness at one layer is fully compatible with concentration at every layer beneath it.

A model can be open at the weights layer and still depend on concentrated semiconductor supply, state-supported training runs no independent party could finance, centralized cloud for inference at useful scale, datasets that cannot be inspected or reproduced, national telecommunications, and alignment and provenance decisions embedded during training that no downstream user can audit. Published weights answer one question—may I run this?—and leave every question in Chapter 17 untouched.

$$\text{open model} \neq \text{open sovereign stack}$$

Open sovereignty additionally requires inspectable or verifiable models, portable execution, hardware diversity, privacy, local or distributed compute, independent energy and communications, user-controlled identity and settlement, and the practical ability to fork or exit. That list is the stack this document specifies, and the gap between it and an open-weights release is the entire subject of Parts III through V.

> A gift at the weights layer is not a transfer of sovereignty. It may be the most efficient way yet devised to make a dependent population feel independent.

The practical instruction follows the same rule the rest of this chapter uses: judge a stack by what its users can do without permission, not by what its license file says.

33.5 The Inversion: Who Is Sovereign

The convergence is structural. The politics are opposite, and the difference is not rhetorical—it appears in specific mechanisms.

Two readings of this table matter more than the rest.

The first is the last row. The open stack's characteristic failure is not oppression but irrelevance: fragmentation, coordination failure, and value that accrues to wrappers rather than to the base asset. That failure mode is the subject of Lemma 13.1 and of Red Lines 6 and 9. Every column has a way of losing, and the thesis is not entitled to compare its best case against another column's worst.

The second is that the table has three columns rather than two, which is a change from how

Dimension	American digital-dollar stack	Industrial-collateral stack	Open bearer stack (this thesis)	Mechanism here
Primary sovereign	The state, via regulated intermediaries	The state	The participant	Item 9
External trust	Reduced by protected domestic capacity	Minimized by state capacity	Minimized by verification	Chapter 17
Internal trust	Concentrated in licensed institutions	Concentrated at the center	Replaced by public checking	VERIFYPRICE
Privacy	Visible to regulators and law enforcement	Visible to authorized institutions	Private by default, selective disclosure	Viewing keys, Layer 5
Identity	Institutional and account-based	Administrative and persistent	Contextual and proof-based	Layer 3, PIDL
Settlement	Tokenized, dollar-denominated, permissioned at the edge	Permissioned, policy-responsive	Non-custodial, bearer-like	Layer 5, PRK
Compute	Strategic national champions	Nationally integrated	Portable, independently verifiable	Layer 4
Redundancy	Allied and reshored supply	Controlled from the center	Produced across independent operators	$\mathcal{O}_s$, Section 17.8
Duration warehouse	Bills, retirement defaults, leveraged intermediaries	Directed banks; state mandate to hold the interval	Not supplied by Layers 0–6; labeled credit beside duration-neutral money (Section 3.3, Section 34.5.1)	Diagnosis, not a layer
Agent operation	Licensed fleets; whether a delegation may leave the perimeter is the operator's decision	State and champion fleets; whether a delegation may leave is the center's decision	Any operator; a hop crosses perimeters on a receipt, if checking stays cheap	Chapter 7, Corollary 30.1
Chief danger	Enclosure by rescue	Administrative enclosure	Fragmentation and weak value capture	Section 5.2, Lemma 13.1

Table 33.1: Three sovereign stacks: convergent structure, different locus of sovereignty. The first two columns differ in method and agree that the state holds the keys.

this comparison is usually drawn. The convenient framing sets a closed foreign stack against an open Western one. The evidence in Section 33.3 does not support it. Two stacks are currently under construction in which the state holds the keys; they differ in method, legitimacy, and the degree of legal recourse available to a participant, and those differences are real and worth a great deal. They do not differ on the question this thesis asks, which is whether a participant may exit without permission.

33.6 What This Validates

Three claims in this document become harder to dismiss as utopian.

Hardened stacks are achievable. A frequent objection to Layer 0 is that verifiable machines, energy provenance, and jurisdictional dispersion are too expensive and too slow to be real. The closed stack demonstrates that vertically coherent physical infrastructure can in fact be built deliberately, at scale, over decades, against market signals. The question is who pays and who benefits, not whether it is possible.

Resilience is priced by serious actors. Section 17.10 argues that redundancy carries option value invisible to normal cost accounting. That argument is not a rationalization invented to defend this stack's overheads. It is the operating assumption of states making the largest infrastructure allocations in the world.

The physical layer is monetary. The thesis's insistence that Layer 0 belongs in a monetary argument, rather than in an appendix about datacenters, matches how sovereigns actually reason about energy: as the substrate of the capacity to honor claims.

33.7 What This Warns

The warning is sharper than the validation, and it is the reason this chapter sits before the objections rather than in an appendix.

> Abundant power and hardened infrastructure do not produce liberty. They produce capacity. Capacity can equip a population or condition it, and the same generating plant serves both.

A stack that is energy-abundant, industrially self-sufficient, computationally sovereign, and administratively integrated has everything required to make participation conditional: electrified transport, domestic payments, integrated identity, state-directed credit, national

cloud, and pervasive telemetry. Nothing in the engineering resists this. The technology is identical; the difference is who holds the keys and who may exit.

This is precisely the participation line of Section 3.6 and the enclosure risk of Section 5.6, arriving not as speculation about a possible Western drift but as a functioning system. Red Lines 11 and 12 exist because the same capacities this thesis wants to build can be assembled into the opposite arrangement, and the difference is not visible from a capability audit.

> The unresolved question of the coming period is not whether systems become more integrated, computational, and hardened. They will. It is whether the people inside them are owners or tenants.

The thesis's normative purpose therefore cannot be "build a hardened stack." That target is already being hit. Its purpose must be to build one whose hardness protects the participant rather than the administrator—which is a claim about key custody, exit rights, disclosure defaults, and value capture, not about throughput.

33.8 What This Chapter Does Not Claim

Discipline here matters, because this material is unusually easy to overread.

Energy consumption is not economic value. Large energy throughput demonstrates industrial scale. It does not establish productivity, capital efficiency, household welfare, or the quality of investment. A system can consume enormous energy while destroying capital, and redundancy financed by suppressed consumption and bad debt transfers costs rather than eliminating them.

Resilience is not automatically superior to efficiency. Excess capacity can become misallocation, local-government liability, environmental damage, and chronically weak returns. The question is never whether resilience is good but whether the insurance was correctly priced—which is exactly why Section 17.8 insists on a measured index rather than a slogan.

The closed stack has not escaped its own constraints. A very large imported-oil dependence remains. Refining flexibility, stockpiles, electrification, and synthetic fuels buy time and bargaining power; they do not create domestic petroleum. The accurate description is increasingly energy-*resilient*, not energy-*sovereign*.

This is not a prediction. This chapter does not forecast which stack prevails, nor asserting that the closed model is ascendant. The claim is narrower and structural: both models are responses to the same collapse of soft guarantees, and they differ in who receives sovereignty.

This is not an endorsement. The strategy is analytically instructive and politically inverted relative to everything in Chapter 4. Describing a system's coherence is not approving of it.

33.9 What This Changes in the Stack

A chapter that only reframed would not belong in Part VI. Four concrete consequences follow.

1. **The competitor is not fiat, and there is more than one of it.** The thesis is usually read as arguing against soft-guarantee fiat and custodial intermediation. The more serious competitor is a hardened, competent, permissioned stack that delivers verification, settlement, and compute with excellent uptime—and conditions access. Section 33.3 indicates two such stacks are being built rather than one, by systems that regard each other as adversaries and that converge on the same answer to the question of who holds the keys. Both win on convenience and capability, not on principle. The thesis must therefore compete on exit rights and value capture rather than on capability alone, and it should expect no jurisdiction to be a natural home.

   Neither closed stack needs to prohibit verification. Each is, among other things, an operator of agent fleets, and a fleet is a trust perimeter (Section 7.2). Inside the perimeter no portable receipt is needed; the operator trusts its instances. The closed stack therefore wins the delegation economy not by banning proofs but by making external checking uncompetitive on convenience: by shipping enough capability that the mesh loses on latency, or by raising verification cost through the means Red Line 13 already names. Corollary 30.1 is the mechanism, and the agent-operation row of Table 33.1 is the question it poses: who operates the agents that issue the delegations, and can a delegation leave the fleet? Identical VERIFYPRICE, identical uptime, opposite answers on whether a hop can cross the perimeter. That is the inversion this chapter already describes, arriving one layer up.
2. **Dispersion is a monetary property, not an operational preference.** If the alternative to a neutral stack is a competent closed one, then jurisdictional and grid dispersion stop being engineering hygiene and become the substance of the neutrality claim. This is what Section 17.8 measures and what Red Line 13 protects.
3. **Capability audits are insufficient.** Two stacks can post identical VERIFYPRICE, uptime, and throughput while differing entirely in whether users can leave. Telemetry must therefore include agency and exit metrics (Section 3.6, Red Line 12), not only performance. A scoreboard that measures only capability cannot distinguish the three stacks compared in Table 33.1—which is the most important thing it could show.
4. **Physical sovereignty requires a financial time horizon.** Energy determines what can be built; duration finance determines whether it remains funded until completion. The closed stack warehouses that interval by command. The American digital-dollar stack currently warehouses it with retirement defaults, market-value indexes, and leveraged intermedi-

aries, or postpones it by issuing bills. The open stack specified here warehouses *claims* that are duration-neutral; it does not, and must not, pretend that Layers 0–6 are a pension. The completeness problem—who finances the reactor—is an objection, not a layer, and is answered in Section 34.5.1.

The missing layer: household transmission. The comparison table above stacks sovereign capabilities against sovereign capabilities, and the stack this thesis defends can lose a comparison the table cannot express. Sovereign resilience and household resilience are different quantities, $R_{\text{state}} \neq R_{\text{household}}$: a state's capacity to sustain energy, industry, defense, and administration can be intact — can even be improving — while the median household's capacity to sustain shelter, savings, family formation, and shock absorption collapses, because the state's continuity was partly financed by consuming the household's buffers (Section 3.6). This is not a redistribution argument; it is a stability argument. A household running a negative operating surplus does not stop consuming — it consumes out of stock:

$$O_i = \underbrace{Y_i + I_i + T_i}_{\text{income}} - \underbrace{E_i + DS_i}_{\text{essentials + debt service}}, \qquad O_i < 0 \Rightarrow \Delta R_i < 0$$

where Y_i is labor income, I_i is investment income (interest, dividends, rent, realized gains actually drawn), T_i is net transfers received, E_i is essential consumption, DS_i is contractual debt service, and R_i is the household's resilience stock: liquid savings, unused borrowing capacity, accessible asset equity, reliable social insurance. The implication $O_i < 0 \Rightarrow \Delta R_i < 0$ is definitional rather than empirical — a negative operating surplus is funded from stock by construction — so the content lies entirely in the threshold applied to the stock. "$\text{Median}(R_i)$ not collapsing" therefore needs a published number, and one is supplied here: the median household's R_i, deflated by the essentials basket E, declines by no more than 10% over the trailing window on which the state model is read, with the cohort series published alongside the median. Aggregate consumption stays smooth the entire way down, which is precisely why the aggregates are a lagging indicator of the exhaustion gradient of Section 3.6 — dollar-weighted measures do not see person-weighted depletion until it surfaces as delinquency, and by then the buffers are gone. The relevance to this chapter's comparison is direct: the closed stack's strongest claim is that it *works*, and "works" measured only at the sovereign layer can mean the state's stability was funded by directing its households down that gradient. **Distributed resilience — R_{state} intact and $\text{Median}(R_i)$ not collapsing, on the -10% real threshold just stated — is therefore a requirement the open stack must be able to evidence, not a social preference appended after the engineering**: a state that survives by transferring fragility downward eventually receives the fragility back through politics, defaults, labor quality, and the demand for control that §3.6 already established as the repression flywheel. The participation-line telemetry is where that evidence lives.

The joint failure mode: services from the closed stack, price from the wrappers. The two adversary chapters of this Part — the closed sovereign stack and the market-realization plane — are usually read separately, and the most likely way this thesis dies is their combination, which no single dashboard currently names. Run the causal chain. The closed stack matures into State 4 service delivery (Section 29.2) and supplies provenance receipts, compliant privacy predicates, and subsidized sovereign compute at quality and price the open stack cannot match for the median commercial workload. That is exactly the bypass channel of Red Line 6 and Condition A of Red Line 14, arriving through a public-sector door rather than a hyperscaler one. Simultaneously, wrappers and custodial products remain the only legally holdable form of the base asset in major jurisdictions, so whatever exposure demand survives is realized entirely through the market machine (Section 13.7). The end state satisfies every element of the thesis's own definition of failure while reading as success on most instruments: the stack's dashboards stay green because the open stack keeps functioning; the price holds or rises because wrapper demand is price demand, not protocol demand; and the base asset becomes a **reference price with no native loop** — quoted, held, and hedged, but never used to transact, prove, or exit.

Three readings of the joint scenario, in order of severity:

- **The red lines already contain it, separately.** RL6 and RL14 Condition A fire on closed-stack service substitution; RL9 fires on wrapper-led realization; RL12 fires on agency use cases failing to materialize while institutional usage grows. What none of them captures is the *joint* condition — all three firing together is the coherent failure, and each firing alone has an innocent explanation (a hyperscaler price war; a bull market; slow enterprise adoption).
- **The diagnostic signature is a triple divergence,** readable before any individual red line trips: Wrapper–Native Growth Gap persistently positive (exposure outgrowing use), convenience-yield telemetry showing no corridor increment as regime pressure rises (Section 26.8.6 — no one pays extra to part with units because the bearer service is not the one being demanded), and Homestead Ratio falling while closed-stack service metrics improve. Any one of these is noise; the triple divergence is the joint scenario arriving. It has a companion reading, not a fourth term: a falling Human-Revocable Share (Section 26.8.9), which says that what native activity remains is agent-controlled with no human root.
- **The thesis should not survive it quietly.** If the joint scenario materializes, the honest reclassification is the one Section 30.6 prescribes for any breach — the asset is a financial product and the stack is infrastructure, and the monetary claim is retired. The purpose of naming the mode here is that the triple divergence is observable *years* before the individual thresholds bind, and a thesis that publishes eighteen red lines should not be killable only by the one failure mode it failed to name.

The household amplifier: why the joint failure arrives non-linearly. The joint scenario above is a demand-side story, and its most likely trigger is an asset repricing that lands on a household sector with no buffers left to absorb it. Before the shock, resilient aggregates are the sum of two individually fragile mechanisms: upper-cohort discretionary spending out of asset gains, and lower-cohort necessity spending out of depletion. An equity or AI-collateral repricing removes the first term; the second cannot expand to replace it, because depleted households have no R_i left to convert into consumption — so aggregate spending does not decline gently, it breaks (Section 33.9). The same structure runs through the sovereign layer: upper-cohort spending carries tax receipts, so the shock propagates into deficits and Treasury stress, and the policy response faces the trilemma Section 29.2 is really about — tighten for inflation, ease for collateral, or build for supply, each worsening another constraint. For this thesis the implication is a measurement one: the base asset's recession-hedging claim is regime-contingent on household buffers, not merely on asset correlations, and the convenience-yield telemetry (Section 26.8.6) is the instrument that would register whether holders actually pay for the holder-side service in exactly that state — the one state where the thesis says it matters.

The birth pathway of State 4, and why it is not the one the table describes. The closed-stack comparison above treats State 4 as a policy choice: a sovereign rationally coordinating energy, industry, compute, identity, and payments into a permissioned whole. There is a second and historically better-attested pathway to the same end-state, which the state does not choose so much as inherit. A financial clearing event in the strategic sectors — the AI build-out's financing structures repriced, as in Section 29.1 — transfers installed infrastructure from weakly financed owners to whoever can absorb it, and in a stress episode the buyers of last resort are sovereign balance sheets, national champions, and directed credit. The state arrives at compute-and-energy coordination not by deciding to build it but by becoming the residual claimant on it; control is a *consequence* of the rescue, not its purpose ([Perez 2002], for the clearing event's transfer mechanism). This matters for the comparison in three ways. It shortens the timeline: State 4 need not be legislated into existence if a single disorderly quarter can deliver it. It changes the binding constraint: a rescue-born closed stack inherits distressed assets and political mandates rather than optimizing ones, so its service quality is a genuine open question rather than a checkbox. And it identifies the exact juncture at which the thesis's two adversaries become one: the market-realization plane's clearing event is the closed stack's acquisition opportunity, and Section 33.9 has already specified the joint signature — the triple divergence of WNG, convenience-yield telemetry, and Homestead Ratio — as the readable warning. The open stack's defense against the crash-born variant is the same as against the legislated one, only more so: exit rights that survive a stress episode, which is Gate 6's entire content.

Chapter 34

Objections & Responses

Any thesis that asks whether a new thing can become money deserves hard pushback, and this one asks it of a stack still under construction.

The objections below are grouped into five blocks. The **monetary** block attacks the central claim and is the one to read if you only read one. The **infrastructure** block disputes whether the stack can be supplied. The **political** block disputes whether power permits it, and whether the project would remain distinguishable from what it opposes. The **market-structure** block disputes whether financialization quietly substitutes for monetary adoption. The **completeness** block asks whether duration-neutral money, even if it works, still leaves civilization unable to finance long-lived assets. Each objection is stated in the strongest form available before it is answered, and several are conceded in part.

34.1 Monetary Objections

These are the objections that matter most, because they attack the claim the thesis actually makes: that a useful capacity can become a store of value. If any of the seven below succeeds, the stack may still be worth building, but it is infrastructure rather than money.

34.1.1 "Why not just Bitcoin?"

Objection: Bitcoin already satisfies the store-of-value requirements this thesis specifies, and has satisfied them on a seventeen-year record under sustained adversarial pressure: credible scarcity on a schedule nobody can change, verification cheap enough to run on a laptop, censorship resistance demonstrated rather than modelled, bearer custody with no issuer, and no duration for policy to pin negative. It achieves this through a work function whose defining property is that **nobody outside the system buys the work**. That is not an inefficiency awaiting correction; it is the source of the objectivity. A puzzle with no external buyer has no customer to lobby, no supplier to subsidize, no regulator with jurisdiction over its demand curve, and no way for anyone to alter what it costs except by spending more of the same resource. Proof-of-useful-work deliberately introduces an external buyer and therefore weakens precisely that property. So everything in this document is either something Bitcoin already

does, or a change that makes the money worse.

Response:

- **The monetary half of this objection is correct, and it is conceded without qualification.** Bitcoin's work function is monetarily superior to proof-of-useful-work, and superior for exactly the reason given. Unforgeable costliness with no external buyer is the strongest available form of the property, because the absence of a buyer is what makes the cost non-negotiable. Introduce a buyer of the work—an enterprise purchasing verified inference, a regulator requiring provenance receipts—and you have introduced a party who can be subsidized, taxed, mandated, prohibited, or captured, and whose demand can therefore be moved by something other than the resource cost of the work. **PoUW trades monetary objectivity for capacity relevance.** That is a trade, and this document is not entitled to call it an upgrade. Any account that describes SHA-256 as waste or heat is wrong, and wrong substantively rather than stylistically: it misidentifies the absent buyer as an inefficiency when the absent buyer is the product.
- **The thesis's claim is narrower than the objection assumes—and narrower than parts of this document have implied.** It is not that a better money is available. It is that Bitcoin does not, by itself, supply portable attestation of computation and provenance, or verified compute. It was never designed to, and there is no defect in its not doing so. Those are separate goods with separate buyers and separate demand curves: a bank's model-risk function buying proof capacity is not in the market for a monetary asset, and would not be satisfied by one. Whether an asset that supplies those goods accrues monetary premium, or merely trades as a priced service, is an open empirical question. The thesis has already conceded most of the ground here: Section 13.3.9 states that fees, burns, and collateral establish a cash-flow claim and a balance-sheet lockup rather than moneyness, and Section 13.3.11 relocates the monetary claim to a convenience yield it declines to size. The honest position is that the question is open and the telemetry exists to answer it, not that the answer is known and favourable.
- **The slots Bitcoin occupies are not the slot this asset can compete for.** Section 13.3.11 works this out and it is not restated here. In summary: gold holds "works when the network does not," Bitcoin holds "works when the state does not," and neither requires reference hardware, a prover market, or bulk power at the moment of use — though Bitcoin still requires network reachability to settle, which gold does not. A triad asset requires power, reachable networks, and obtainable reference hardware in order to deliver anything at all, which confines it to a narrower band—adversarial conditions severe enough to need proofs and privacy, but not severe enough to take the grid. Part I argues that band is the likely one. Arguing it is likely is not the same as sizing it, and the thesis does not size it.
- **In this architecture the thesis is a consumer of Bitcoin, not a competitor to it.** Every privacy corridor specified at Layer 5 has Bitcoin on one leg — BTC↔XMR deployed, BTC↔shielded-

ZEC gated (Section 23.3); VERIFYSETTLE is defined over those corridors (Section 23.5); the deployed non-custodial corridor is a BTC↔XMR atomic swap with Bitcoin on one leg, classified as a corridor rather than a bridge (Section 23.4); federated e-cash mints extend Bitcoin's privacy at community scale through federated custody — a Class B/C pattern (Section 23.4) that the thesis lists but does not count as non-custodial settlement (Section 29.3); and receipt ledgers anchor to a neutral chain, of which Bitcoin is the obvious candidate. Bitcoin is also named as a bridge asset that survives the transition rather than a precursor to be discarded (Section 29.1.1). A document that wanted Bitcoin displaced would have to explain why it built its settlement layer on top of it. Damage to Bitcoin is damage to this stack.

- **What losing to Bitcoin would look like, specified so it can be checked.** The thesis loses this argument if triad capacity is consumed at scale—receipts, verified units, private settlement all growing, the capturable wedge intact, native fees tracking usage—while the asset shows no behaviour distinguishing it from a claim on a service business, and the monetary bid accrues to Bitcoin instead, across repeated episodes of the regime pressure the hedge is supposed to be for. That pattern is named and its measurement specified as **Service-Good Realization** (Section 30.7). It remains a market-realization warning because its price-comparative clause cannot be a red line without breaking Corollary 13.1. Red Line 15 instead covers the protocol-observable native collateral–capacity spiral.

The summary belongs here before any critic states it: **Bitcoin is the better money; the triad supplies different goods, and whether those goods carry monetary premium is unsettled.** That is a weaker claim than a succession narrative and a considerably more defensible one, and it is what Section 13.3.11 already implies. Readers who came here expecting the thesis to argue that Bitcoin is superseded should note that no such argument is made anywhere in this document.

34.1.2 "The buyer concession is repaired by agents."

Objection (the friendly version, which is the more dangerous one): The previous response concedes that useful work has a buyer and a buyer has a jurisdiction. But Section 7.4 says the delegation buyer is a mesh with no procurement desk. A mesh cannot be mandated, subsidized, or withdrawn as a whole. So the concession to Bitcoin can be dropped, proof-of-useful-work no longer trades objectivity for relevance, and the thesis should say so.

Response:

- **No.** The concession stands and this document does not withdraw it. The reason is in the concession's own text: a puzzle with no buyer has no demand curve anyone can move. A mesh of independent agents still has a demand curve; it is moved operator by operator rather than all at once. Harder to coerce is not the same as impossible to coerce, and Bitcoin's work function has the second property.
- **What the delegation model changes is the shape of the coercibility, and it names the**

instrument. Fleet-originated demand is coercible at one desk. Independently originated demand is coercible at as many desks as there are operators. The agent-originated demand share by operator class is the series that says which the network has, and the Bitcoin comparison should be run against that series rather than against the thesis's hope.

- **The limit case is worth naming because it is not reachable.** As the operator count goes to infinity and each operator's share goes to zero, the delegation buyer approaches an absent buyer. It never arrives, because every operator has a jurisdiction and a power bill, and Red Line 13 is where that fact reads. A thesis that claimed the limit had been reached would be claiming that the agent economy is already a mesh, which is the empirical question Corollary 30.1 leaves to $r(W)$.

The summary, stated here before a critic states it: **the delegation buyer is the least coercible external buyer a useful-work function can have, and an external buyer is still the trade Bitcoin declined to make.**

34.1.3 "This is just another chain / coin."

Objection: We've seen this movie: every new system claims to be hard money with a fancy narrative. This is just buzzwords on top of a token.

Response:

- The core claim is not "this token pumps," but "Privacy, Proofs, and Compute are verifiable necessities that can be priced and audited."
- The focal metrics are **VERIFYPRICE/Reach/Settle and triad usage**, not just TVL or number of wallets.
- Work Credits are claims on **measurable work**: proofs, private settlements, verified FLOPs, each extinguished only by verified delivery, with energy and hardware profiles tied in.

If, in practice, the system behaves like "just another chain" (no canonical workloads, no PoUW, no privacy corridors, no telemetry), then the objection is correct. The point of the architecture is to make that failure **obvious**, not to hide it.

The objection has an empirical record, and the thesis should cite it against itself. "Utility does not imply money" is not a conceptual puzzle; it is a decade of measured negative results. Filecoin, Render, Akash, Golem, and Livepeer each bootstrapped real infrastructure, real usage, and real fee-bearing demand — storage, GPU rendering, compute, and video transcoding respectively — and in none of them did native fees approach a meaningful share of the security budget, and in none did the asset acquire monetary behavior beyond beta to the sector. Fee coverage on the closest precedents has remained in the low single digits of issuance for years. (Data caveat: those fee-coverage figures are read from public analytics dashboards — Messari [Messari 2024–26] and Token Terminal [Token Terminal 2024–26] protocol revenue series against reported issuance — not from audited financials, and the ratio is sensitive to how

each dashboard classifies fees, burns, and incentive emissions. The order of magnitude is not in dispute; the second digit is.) This is the strongest available form of the objection, and it deserves the strongest form of the answer:

- **The precedents did not satisfy the conditions, and the thesis does not dispute the outcome.** Each of those systems fails or partially fails the Value Capture Lemma conditions as this thesis specifies them — fees payable in alternatives, minimal burns, weak collateral lockup, and (most decisively) service demand largely bypassable through centralized equivalents at similar quality. Their observed failure of premium accrual is what the lemma predicts, not a counterexample to it.
- **But the precedent cuts the other way too.** The II→III phase gate (Section 29.3) demands fee coverage of $\geq 30\%$ of the security budget sustained for 12–24 months — roughly an order of magnitude above anything the closest precedents achieved. The honest reading is not that the gate is modest but that it is deliberately set where no comparable system has reached, because below it the security budget is issuance and the asset is, economically, exactly what the objection says: a token on top of infrastructure. The gate should be read as the thesis betting against the precedent set, and Red Line 5 exists to settle that bet in public rather than around it.
- **What would distinguish this attempt.** Not narrative, and not telemetry theater: fee-bearing demand that cannot route around the asset (Condition 5, and see Section 13.5.1 for the incumbent channel), verified at the workload tier (Section 22.10). Until the boards show it, the default expectation informed by the precedents is that this is another service token, and the burden sits on the design, not on the reader's imagination.

34.1.4 "The triad does not belong in one token."

Objection: Privacy settlement, proof procurement, and verified compute have different users, bottlenecks, collateral requirements, and political risks. Combining them in one token imports unrelated failures, cross-subsidizes weak services, and mistakes shared branding for monetary coherence.

Response:

- The objection is valid against any architecture that assumes one-token coherence. This thesis does not assume it.
- The Triad Coherence Test (Section 8.3) compares one native monetary asset, a neutral reserve plus service-specific credits, and shared settlement plus modular domain collateral.
- Architecture A survives only if common security and liquidity create more value than cross-domain governance and wrong-way risk destroy. If Architecture B or C performs better, the thesis narrows to an open service stack beside an existing reserve asset.

34.1.5 "Utility does not imply money."

Objection: Electricity, bandwidth, and cloud compute are all indispensable, but claims on them are not stores of value. Why should Privacy, Proofs, and Compute be different?

Response:

- The objection is correct as stated. Utility alone does not create money. The Value Capture Lemma establishes only whether service value accrues to the base asset through fees, burns, collateral, issuance discipline, and constrained bypass.
- Monetary treatment additionally requires stress-deliverable service, a persistent self-custodied loss-bearing holder constituency, and the holder-side service flow of Section 13.3.11. If those fail, the system can be useful infrastructure and the base asset a service or cash-flow claim without being a store of value.
- The telemetry regime makes the distinction testable: the Value Capture Board reads service accrual, the Native Monetary Buyer Map reads risk absorption, and the Layer 0 board reads DVC.

34.1.6 "The asset can be bypassed."

Objection: Users may need proofs, privacy, and compute, but they can buy them from AWS, a prover marketplace, a privacy wallet, or a stablecoin-based service without holding the native asset.

Response:

- The objection is correct unless the protocol enforces native value capture. The SoV thesis requires that core fees, collateral, slashing, priority access, settlement, or governance-critical operations be denominated in the native asset; that a material share of fees be burned or retired; and that operators lock the asset to provide service.
- If users can consume equivalent triad capacity without touching the asset, the system may be useful infrastructure, but the monetary thesis fails. This is Red Line 6 (Section 30.6).
- The honest answer is: this risk is real and must be continuously monitored. The Economic Coverage Board and Value Capture Board (Chapter 26) track native-asset fee share and bypass indicators.

34.1.7 "Gold still beats this because gold needs no network."

Objection: Gold requires no software, no telemetry, no cryptographic agility, and no functioning power grid. Any network-dependent asset is structurally inferior as a store of value.

Response:

- Correct under some failure modes. That is why gold is a bridge and permanent non-digital redundancy asset (Section 29.1, Section 29.1.1), not a discarded precursor.
- The triad is not "better than gold" universally; it is better suited to dense digital civilization

if the network remains reachable and verifiable. The two are complements across different failure surfaces, not substitutes competing for the same use case.

34.2 Technical and Infrastructure Objections

The next group disputes whether the triad can be supplied at all under adversarial conditions—whether anyone wants it, whether useful-work markets can stay decentralized, whether proofs deliver what they appear to, and whether the physical substrate can be denied to it.

34.2.1 "Users don't care about privacy or proofs."

Objection: People trade convenience for privacy all the time; most don't verify anything. Why build a SoV thesis on properties most users won't touch?

Response:

- Users may not **ask** for privacy or proofs, but they **suffer** when they are absent: identity theft, surveillance, misinformation, censorship, frozen funds.
- The stack is not asking end-users to run verifiers or design circuits; it is asking:
  - infrastructure operators to run verifiers,
  - institutions to demand receipts, and
  - developers to call PAL/PRK instead of opaque APIs.
- The "user" that cares most may be a treasury, DAO, insurer, or regulator—actors who must explain themselves.
- The objection also assumes the user is a person. The marginal user of proofs in a delegation economy (Chapter 7) is a process that cannot do consequential work across a trust perimeter without them: it has no way to check a stranger's work except by verifying it, no exit except refund safety, and no basis for accepting a unit except that anyone could check it. Whether humans come to care about proofs is a question about preferences; whether agents require them is a question about mechanism, and the second answer does not depend on the first. This does not make the demand monetary—Red Lines 16 through 18 still decide that—and it does not make the demand open, since a process inside one operator's fleet needs no portable receipt at all. It does mean that the constituency the objection says does not exist may arrive before the one it is arguing about.

If, after a decade of increasing repression and AI-mediated reality, nobody is willing to pay for privacy, proofs, or verified compute, then the triad does not become money. The thesis is that the opposite is more likely.

34.2.2 "PoUW will centralize / can't compete with hyperscalers."

Objection: Useful-work mining sounds good, but hyperscalers and incumbents will always dominate; you'll just rebuild cloud oligopolies inside a blockchain.

Response:

- Hyperscalers already dominate raw compute; the point of PoUW is not to out-compete them on price, but to:
  - turn **verified units of useful work** into a commodity;
  - ensure **anyone** can verify;
  - keep entry for provers open at the margin.
- Layer-0/4/6 design fights centralization by diversifying hardware profiles; measuring and publishing prover concentration; using rewards and Work Credit policies to favor diversity; allowing hyperscalers to participate, but not to be **the only ones**.

If the market decides that centralized compute is always "good enough" and verifiability never matters, then AI Money stays a nice phrase. The thesis is that high-stakes actors—finance, defense, safety-critical infrastructure—will demand **verifiable** compute from multiple vendors.

34.2.3 "Proofs do not prove truth."

Objection: The thesis calls these instruments "Attestation Money" and speaks of proofs as if they settle disputes. But proofs only verify computation under stated assumptions. They don't tell us whether the inputs were honest, the model was appropriate, or the conclusion is socially meaningful.

Response:

- Correct. This is why the thesis uses "attestation" rather than "truth." Proofs do not create truth; they bound disputes. They make specific claims cheap to verify: origin, custody, computation, policy compliance, and settlement finality.
- The thesis is not that cryptography tells us what is true, but that it reduces the surface area over which institutions must be trusted. That reduction is valuable even if it is bounded.
- If the thesis overclaims—if it suggests that proofs solve epistemology rather than computation verification—the reader should substitute "bounded attestation" wherever "truth" appears. The economic argument does not require semantic truth; it requires cheap, public verification of specific claims.

34.2.4 "They'll just shut off the internet / app stores."

Objection: Sovereigns can simply block the network or remove apps from stores.

Response:

- Total, permanent shutdowns are blunt and politically costly; partial, targeted throttling is

more likely. The stack assumes this and designs for it:

- Multiple obfuscated transports (Layer 1).
- Content-addressed updates and offline installers (Layer 2).
- Sideload paths for clients.

- Treat reachability and update-health as SLOs (VerifyReach): if clients in censored regions can still fetch proofs and updates via at least one path, comms resistance is doing its job.

34.2.5 "This burns too much energy."

Objection: PoUW is just PoW with extra steps; it still wastes energy.

Response:

- All monetary substrates consume scarce resources—geology, enforcement, balance-sheet capacity, or energy. The premise of the objection is that SHA-256 hashing is waste, and that premise is rejected here as firmly as the objection is: the energy buys unforgeable costliness under a work function nobody outside the system can price, and that is what the money is made of (Section 34.1.1). PoUW does not correct a defect. It binds the same class of expenditure to work that has an independent buyer.
- The right question is energy per verified FLOP or per proof unit, and whether that trend is improving. If energy-per-receipt falls while verified-capacity-per-token rises, the system is delivering more *verified capacity* per unit of energy. It is not thereby delivering more *monetary objectivity*: on that axis a work function with an external buyer is weaker, not stronger, and the comparison should not be run in a single direction and reported as a win.
- **Measurable answer:** Facility Energy Receipts (FERs) and "work per kWh" metrics make energy usage auditable. Energy efficiency shows up in VERIFYPRICE (cost component) and in Layer-0 telemetry.

34.2.6 "You will lose the competition for electrons."

Objection: This is the serious version of the energy objection, and it is not about waste. Verification workloads will compete for power against industrial demand and state-backed AI build-outs. Those competitors have sovereign balance sheets, national-champion status, and priority in the interconnection queue. Third-party proving loses that contest, and no amount of telemetry changes the outcome.

Response:

- **This objection is substantially correct**, and the thesis now treats it as a threat class rather than a talking point (Section 5.5). Curtailment, tariff discrimination, interconnection denial, and load prioritization are real instruments, administratively cheap, and rarely legible as repression.
- **But it applies to proving, not to verification.** These are different exposures (Section 17.10.1).

Verification runs on reference hardware at the edge in small amounts; its exposure is hardware obtainability, not bulk power. Losing the bulk-power contest raises the cost of *producing* proofs. It does not by itself break "anyone can verify," which is the hinge.

- **Where it does bite, the thesis measures it and lets it bind.** Sovereign optionality (Section 17.8) aggregates curtailment exposure, fuel mix, backup depth, and jurisdictional risk into a published index; risk haircuts price fragility into issuance (Section 25.6.5); Red Line 13 retires the thesis if verification affordability becomes a sovereign policy variable.
- **The honest limit:** if a jurisdiction decides third-party proving will not be energized, cryptography does not overrule the grid. The response is dispersion and honest reporting. A thesis that claimed otherwise would be lying about physics.

34.3 State, Political, and Governance Objections

These objections concern power: whether states will tolerate lawful privacy, whether the project's own instruments can be turned against their users, whether a state-integrated stack simply outperforms an open one, and whether the thesis is analysis or preference. The most serious charge in this group is not that the project loses, but that it wins and becomes the thing it objected to.

34.3.1 "Governments will never allow lawful privacy at scale."

Objection: Sovereigns will not tolerate unbreakable privacy and bearer-like digital money; they will regulate and block until these systems are marginal.

Response:

- Some governments will indeed oppose; others will see advantages in verifiable, receipt-backed compliance; lower settlement risk; reduced dependence on foreign platforms and currencies.
- The architecture is about **resilience**, not legal victory: it assumes partial repression; spreads hardware, comms, and governance across jurisdictions; keeps protocol-level custody and identity neutral.
- Lawful privacy is engineered, not begged for: viewing keys and receipts allow compliance without re-centralizing custody.
- **Critically:** The system designs *neutral infrastructure with consented disclosure*, not "dark finance." Policy compliance is predicate-based (ZK proofs of allowlist membership, authorization, jurisdiction), not graph-inspection.

If all major jurisdictions converge on banning any form of non-custodial digital value, a lot more breaks than this stack. The design goal is: **if even a few open jurisdictions remain, and some gray-market paths exist, can the triad continue to function?**

34.3.2 "Receipts will become surveillance tools."

Objection: All these receipts and proofs will just become a new surveillance layer. Governments will demand access; the system will comply; privacy dies.

Response:

- Receipts are **privacy-preserving by design**: they prove predicates (membership, compliance, authorization), not identities. A receipt proves "sender is in cleared set" without revealing which entry.
- Disclosure is **consented and scoped**: viewing keys reveal specific flows to specific auditors for specific time windows, not universal access. There is no "master key."
- **Constitutional constraint:** "Policy = predicates, not graph inspection" is a hard constraint. Any policy requiring universal tracing is treated as incompatible with the monetary design.
- The system **explicitly rejects** policies that require global traceability. If a regulator demands "show me all transactions," the answer is: "We can prove compliance with specific predicates; we cannot provide a surveillance feed, because the architecture doesn't support it."

If receipts become surveillance tools, the design has failed. The telemetry should make this visible: if most flows require real-name disclosure, the Settlement & Privacy Board will show it.

34.3.3 "Permanent receipts make obedience rational."

Objection: A stack that makes every decision reconstructable changes behavior before the decision is made. When an agent's standing instruction is followed and the result is bad, the loss is charged to the instruction; when a person overrides it and the result is bad, the loss is charged to the person, and the receipts preserve the warning, the departure, and the loss with perfect fidelity. The rational response is never to override. A stack built to preserve human agency ends by teaching its principals to hand over their judgment along with the keys.

Response:

- **The mechanism is real, and the thesis names it as a failure path.** Standing held on paper and obedience in practice is standing inversion without the key failure (Section 5.2). It is not dismissed here.
- **The record this stack keeps is not a record anyone can read.** Receipts prove predicates; disclosure is scoped and at the holder's discretion (Section 23.7); and lawful privacy is argued throughout as the technical form of the fresh start (Section 3.6). The reconstruction the objection fears requires a reader, and by default there is none. A good-faith departure that turned out badly is a fact its holder may keep.
- **Where a reader is compelled, the stack does not help, and it does not pretend to.** Inside an audited institution a regulator or a court can require disclosure, and there the stack makes reconstruction cheaper while doing nothing about how the institution charges the cost of an

override. That allocation is governance, not protocol. The thesis's own governance holds itself to the standard it recommends: an override record is read across many decisions to judge the rule, not against one outcome to convict the signer (Section 25.5.2).

- **What would show the objection winning:** a human-revocable share that holds while the policy concentration of the same cohort rises (Section 26.8.9)—every root held by a human, none of them used. That reading is on a board.

34.3.4 "This just gives the state better compliance tools."

Objection: Selective disclosure, receipts, and viewing keys sound like privacy, but they are really a more efficient surveillance and compliance apparatus for the administrative state.

Response:

- The stack must distinguish selective disclosure from global graph inspection. Predicates ("sender is in cleared set," "amount is within range") reveal facts, not biographies. If proofs become dossier infrastructure, agency preservation fails—this is the ninth SoV requirement (Item 9) and is tracked directly by the social coercion metrics (Section 27.1.2) and the Agency Preservation Board (Section 26.9).
- The constitutional constraint in Section 27.1 ("policy = predicates, not graph inspection") is a hard design boundary, not a marketing claim. Any policy requiring universal traceability is treated as incompatible with the monetary design.

34.3.5 "The closed sovereign stack simply wins."

Objection: Chapter 33 concedes the case against itself. A state that controls energy, industry, compute, payments, and identity will out-build any open network, and will deliver better uptime and lower cost while doing it. Open verifiable systems are a luxury good for jurisdictions that have not yet had to get serious.

Response:

- **On capability, this may be right**, and the thesis does not contest it. A competent closed stack can plausibly beat an open one on cost, latency, uptime, and build speed. Directed capital does not wait for market clearing.
- **Capability was never the claim.** The thesis argues that certain properties—non-custodial settlement, private-by-default disclosure, portable identity, independently checkable receipts, practical exit—are what make an asset a store of value under repression (Chapter 4). A closed stack does not provide those at any level of capability, because withholding them is what makes it closed.
- **The two are not distinguishable by performance metrics**, which is exactly why the scoreboard must include agency and exit (Section 33.9). Identical VERIFYPRICE and uptime, opposite answers on whether a user can leave.

- **The real risk is not losing to it, but becoming it.** Red Lines 11 and 12 exist for this. An open stack that ends up permissioned, hyperscaler-hosted, and disclosure-mandated has not been outcompeted; it has converted. That is the failure mode worth monitoring, and it is monitored.

34.3.6 "A perfect ledger hardens an unequal starting state."

Objection: Cryptographic verification solves alteration, double-spending, and provenance — then faithfully preserves whatever distribution existed when the snapshot was taken. If incumbents arrived with the assets and entrants must acquire them from wages at post-repricing valuations, a trustless ledger does not widen access; it makes the incumbents' claims cheaper to hold and harder to contest. The stack could win technically and reproduce, more efficiently than fiat did, exactly the exhaustion dynamics of Section 3.6.

Response:

- **The objection is correct about ledgers and has always been priced in.** Nothing in the thesis claims verification redistributes. The Value Capture Lemma's conditions govern whether service value accrues to the asset; they are silent on who held the asset first, and a perfectly neutral ledger is neutral about incumbency too. Conceding this is not a retreat — the thesis never claimed otherwise.
- **But neutrality of record is not neutrality of access.** The fiat comparison is not distribution-neutral either: it is a system in which the mapping from discipline to advancement (Section 3.6) failed while every instrument read green. The relevant comparison for the open stack is whether it *broadens the acquisition channel* relative to that baseline — open-admission proving markets, fee-denominated entry at small denominations, collateral that is capacity-earned rather than capital-inherited, no institutional gatekeeper between a wage and a claim. These are architecture choices, they are already in the design (open admission, Section 34.2.1; Tier C service markets), and they are what distinguishes "neutral about incumbency" from "rigged for incumbency."
- **The stack must be measured on the capability layer, not the formal layer.** Formal access (a valid balance) can coexist with practical exclusion (no plausible path from wages to a meaningful claim). The distinction and its telemetry belong to the ninth requirement (Item 9) and the participation line; the Agency Preservation Board and the buyer-map's cohort telemetry are where a widening acquisition gap would surface. A stack whose self-custodied cohort stops growing from income and grows only from inheritance has failed a requirement even if every SLO holds.
- **What would refute the response:** if, at material adoption, new-participant acquisition of the base asset from labor income stays negligible while incumbent concentration rises — visible in buyer-map cohorts and holder concentration — then the open stack's acquisition channel claim is false and the objection stands in its strong form. That is a measurement,

not a value judgment, and it is on a board.

34.3.7 "This is political preference dressed as analysis."

Objection: If both stacks minimize trust and both are hardened, the preference for the open one is ideology. A national stack delivering energy, payments, healthcare, and compute serves more people better than a fragmented network of cryptographic hobbyists. Calling one "sovereignty for the person" is branding.

Response:

- **The normative commitment is real and stated**, not smuggled: the thesis's purpose is preserving practical agency under monetary instability and administrative integration (Item 9). It does not pretend to be value-free, and a document that claimed neutrality here would be less honest, not more.
- **But the analytical content is separable and falsifiable.** The claim is conditional: *if* soft guarantees weaken, *then* a bearer asset may earn premium only where service remains deliverable, value accrues without bypass, and holders bear loss—measurable through DVC, fees, burns, collateral, settlement health, buyer quality, and verification cost. Section 30.6 lists eighteen ways to show it false. Ideology does not usually publish its own kill conditions.
- **The competence objection is conceded.** A closed stack may well deliver more material welfare to more people. The thesis's response is not that this is untrue but that it is a different question from what happens to holders when the administrator's incentives change and exit is unavailable.
- **What would actually refute the thesis** is set out in Section 30.6: if institutions restore broad credibility without increasing administrative control, if users retain practical exit despite integration, and if open proof systems fail to provide economical alternatives, the thesis is wrong regardless of anyone's preferences.

34.3.8 "AI will be enclosed anyway."

Objection: Hyperscalers have too much capital, data, and distribution advantage. Decentralized verified compute will never be more than a rounding error next to the walled gardens.

Response:

- The thesis does not require decentralized verified compute to be cheaper than hyperscalers (Section 34.2.2). It requires it to be available, verifiable, and censorship-resistant when those properties carry a premium—the same logic developed for AI Homestead vs. AI Enclosure in Section 5.6.
- The Homestead Ratio and Enclosure Risk Flag make the concentration question falsifiable rather than rhetorical. If hyperscaler share of VERIFYPRICE-tracked capacity stays persis-

tently above threshold, Red Line 11 (Section 30.6) triggers and the thesis concedes the enclosure path is winning.

34.3.9 "Agents will be enclosed with the compute."

Objection: The delegation economy is a hyperscaler product. The agents that issue delegations run on hyperscaler models, inside hyperscaler runtimes, on hyperscaler hardware, under hyperscaler terms of service. A hop that begins and ends inside one fleet needs no portable receipt, and most hops will. Cross-Perimeter Share will be a rounding error, agent-originated demand from independent operators will be near zero, and Chapter 7 has described a demand that enclosure has already captured.

Response:

- **On current volume, the objection is probably correct**, and Section 7.4 says so: fleet-originated demand is where early hop volume should be presumed to sit, and the part of it that never leaves the fleet is the bypass channel Red Line 6 already names. Nothing in the delegation model requires that fleets be small.
- **The claim is about the margin, not the mass.** A fleet is a trust perimeter, and the reason a hop stays inside one is that crossing costs more than it returns. That cost is set by $r(W)$ and by the availability of settlement a stranger can rely on. Where those fall, hops can cross; where they do not, hops stay. The model does not predict which; it names the variable (Corollary 30.1) and the instrument (CPS).
- **Open weights do not settle this, in either direction.** Section 33.4 already warns that open weights are not open sovereignty. The same discipline applies here: an open-weight agent running on rented hyperscaler hardware under a real-name account is a fleet instance with a public license file. What makes a hop independently originated is the operator's ability to obtain power, hardware, reach, and settlement without permission, which is the stack the thesis specifies and not the model's license.
- **What losing looks like, so it can be checked.** CPS flat or falling across four quarters while VERIFYPRICE stays inside SLO and receipt volume grows. That is enclosure by convenience: the receipts are being produced and consumed inside perimeters that do not need them to be portable. It is the leading indicator of Red Line 11 and should be read before capacity concentration, not after. Until receipts carry the field that makes CPS computable, this check cannot be run, and the thesis says so rather than reading the silence as good news.

34.3.10 "Governance will just re-centralize."

Objection: The stack is too complex; whoever runs it will become the new chokepoint.

Response:

- Complexity does not have to mean opaqueness. The governance layer is deliberately thin:

decisions are about parameters and SLOs, not about picking winners.

- Multi-jurisdictional foundations, transparent config changes, and slashing rules keyed to public telemetry (VERIFYPRICE, decentralization stats, corridor health) make governance legible.
- If a network cannot show **who changed what, when, and in response to which metrics**, it is not an SoV candidate—however elegant its whitepaper.

34.4 Market-Structure and Co-option Objections

The fourth group is specific to the Market Realization Plane, and it is the group most likely to be dismissed by protocol designers as somebody else's problem. It concerns what happens when the asset succeeds financially without succeeding monetarily—and in particular the tempting inference that a rising price settles the argument.

34.4.1 "The asset will be co-opted, not bypassed."

Objection: Even if the protocol enforces native value capture, the financial system will co-opt the asset through ETFs, treasury vehicles, margin loans, and regulated custody. Most holders will never use self-custody, privacy rails, or verification tools. The asset becomes a speculative vehicle inside the existing financial system, not an independent monetary base.

Response:

- The objection describes a real and ongoing dynamic. The thesis acknowledges it as a variant of bypass risk: co-option routes demand through custodial instruments that do not exercise the fee-burn-collateral loop.
- The defense is protocol-level: if privacy, settlement, and proof capacity can only be accessed through native protocol interactions (not custodial wrappers), then custodial holders have exposure to price but not to the monetary properties that justify the price. Over time, this creates a divergence: custodial exposure without utility is speculation; protocol-native usage with utility is monetary demand.
- The telemetry regime must track this: native-asset fee share vs. custodial wrapper volume, protocol-native settlement vs. exchange-settled trades, and self-custody ratio vs. custodial concentration. If custodial wrappers dominate and protocol-native usage stagnates, the co-option objection wins.
- The strongest counter-argument is that the protocol stack must make self-custody and protocol-native usage *easier and more rewarding* than custodial alternatives—not by ideology, but by delivering privacy, verification, and settlement that custodians cannot replicate.

34.4.2 "This will just become another ETF."

Objection: Regulated wrappers, treasury companies, and ETFs will absorb most institutional demand. The asset will become a reference price for a financial product, not a monetary rail.

Response:

- It might. That is why the Wrapper Dominance Ratio exists (Section 13.6). Price exposure without protocol usage is not monetary adoption, and the thesis names this failure mode explicitly rather than hiding behind price action.
- The telemetry regime is designed to catch this early: if WDR rises while native fee share, private settlement, and collateral lockups stagnate, Red Line 9 (Section 30.6) triggers before the price signal would suggest anything is wrong.

34.4.3 "The price went up, so the monetary thesis is working."

Objection: The asset has appreciated substantially and institutions are buying it. That settles the argument.

Response:

- It settles nothing. Price appreciation can be produced by return-decoupled wrapper allocations, leveraged rebalancing, dealer hedging, passive inclusion, treasury-company issuance, or trend following — none of which touch a fee, a burn, a proof, or a settlement (Corollary 13.1).
- Monetary adoption requires native fee demand, burns, collateral, private settlement, proof consumption, verified compute, and healthy stack telemetry. Price is neither necessary nor sufficient evidence for any of these.
- The symmetric claim also holds: a falling price is not falsification. The thesis is tested on the Value Capture Board (Section 26.9), and interpreted — never tested — through VerifyFlow (Section 26.8).

34.4.4 "Wrappers increase access, so co-option is harmless."

Objection: ETFs and custodians expand the buyer base, deepen liquidity, and give institutions a legal path in. That is unambiguously good for the asset.

Response:

- Partly true, which is what makes it dangerous. Wrappers can improve liquidity and institutional holdability. They can also centralize custody, bypass native value capture, reduce self-custody, and manufacture holders who never exercise the monetary properties being capitalized (Section 5.6).
- The outcome is empirical, not ideological. The question is whether wrapper growth *seeds* native use or *substitutes* for it, and the Wrapper–Native Growth Gap answers it directly (Section 26.8).

- Note what this thesis does *not* claim: that every ETF is an adversary. Unlevered wrappers can be genuinely useful. The failure is not their existence; it is the inability to distinguish accessibility from adoption.

34.4.5 "Market mechanics are external, so protocol designers can ignore them."

Objection: Wrappers, dealers, and fund flows are outside the protocol boundary. Designers should build the stack and let markets do whatever markets do.

Response:

- External wrappers reach back inside. They affect treasury behavior, collateral demand, governance concentration, liquidity-provider incentives, public understanding, security-budget expectations, regulatory pressure, and the price that the protocol's own participants use to make decisions.
- The correct posture is narrow but firm: the protocol must not *manage* its price, and must nonetheless *instrument* the structures through which its monetary claims are represented (Section 13.7). Measurement is not intervention.
- Ignoring the plane does not make it inert. It makes it unmeasured, which is precisely the condition under which the compositional adversary (Section 5.2) does its work.

34.5 Completeness Objections

This last group is not an attack on moneyness. It is an attack on the claim that moneyness is enough. A duration-neutral, verifiable, bearer-capable asset can still leave the physical stack unbuilt.

34.5.1 "Who finances the reactor?"

Objection: Neutral money and verifiable settlement can preserve value. They cannot, by themselves, finance a nuclear plant whose construction takes seven to fifteen years depending on jurisdiction and whose productive life spans sixty. Gold has no coupon. Bitcoin has no underwriting desk. Proofs establish that a computation occurred. Privacy prevents informational capture. None of these converts present savings into long-lived productive capacity. AfterFiat has told savers to exit the bond, then left civilization with no way to build the plant. That is not a monetary design. It is a savings technology with a stack diagram.

Response:

- **The objection is correct as a completeness claim, and it is conceded as one.** Duration-neutrality of the monetary object (Chapter 4) is a requirement of repression resistance. It is not a theory of capital formation. Section 3.3 distinguishes duration of the claim from duration of the project for exactly this reason. A post-fiat order that refuses credit altogether

will not build Layer 0. A post-fiat order that cannot tell collateral from credit will recreate the opaque leverage and maturity mismatch this thesis is trying to escape.

- **The thesis is not the duration warehouse, and Layers 0–6 are not a pension.** Reserve collateral should remain duration-neutral, politically hard to pin, and bearer-capable. Duration finance should remain labeled credit: underwriting, covenants, maturity transformation, loss-bearing equity, and institutions that can sit through marks. Proofs can make construction progress, energy receipts, Facility Capacity Receipts, covenants, and restructuring states checkable. They cannot delete time risk, inflation risk, construction risk, political risk, or technological obsolescence. If a project cannot attract a loss-bearing holder of time, it does not get built—and that failure is not a monetary-design bug.
- **Treasuries currently fuse the two functions.** They are used as safe collateral, as a yield curve, as pension-matching assets, as repo, and as the mechanism that converts savings into public expenditure. Section 3.3 shows the duration function weakening while the collateral function is still treated as intact. Replacing Treasury reserve demand with gold or Bitcoin would not automatically replace duration and credit. The world would still need instruments through which savers finance governments, infrastructure, industry, housing, and corporate investment. The design problem AfterFiat actually poses is how to *separate* those functions without recreating discretionary bailout structures.
- **China's answer is command; the current American answer is a rule.** Directed banks can be compelled to warehouse the interval. Default enrollment into market-value-weighted bond indexes cannot. Neither is the open-stack answer. The open-stack answer is: money stays duration-neutral; credit stays labeled and auditable; governance decides who may write the default and who bears the loss when the rule fails (Chapter 25, Section 3.6).
- **What would make the objection fatal rather than clarifying.** If the thesis began treating Work Credits, proofs, or the native asset as a substitute for long-duration project finance—a coupon by another name, a construction bond wearing a monetary label—then duration-neutrality would have been undone and the bondholder's trap would have been re-entered through the protocol door. That is a use this document refuses.

34.5.2 The Open Duration Warehouse

"Credit must remain labeled credit" is a boundary, not a financing architecture. An open stack therefore permits explicit duration-bearing instruments beside the duration-neutral base asset:

- proof-audited project notes with stated maturity, covenants, default states, and loss priority;
- revenue-linked capacity bonds whose payments depend on attested delivered service rather than promised nameplate output;
- power-purchase and compute-purchase agreements that contract defined quantities, locations, delivery windows, and force-majeure states;
- senior, subordinated, and first-loss infrastructure tranches with disclosed attachment and

exhaustion points;

- pre-purchased capacity contracts that transfer delivery and basis risk explicitly;
- separately capitalized first-loss or insurance facilities with published exclusions and claims waterfalls;
- bankruptcy-remote project entities that isolate construction and operating claims from the base-asset constitution; and
- continuously attested cash-flow, collateral, construction-milestone, and restructuring waterfalls.

Proofs can reduce information asymmetry, attest milestones, audit revenue, and make covenant breaches public. They cannot eliminate default, commodity risk, construction delay, technology obsolescence, political seizure, or the need for loss-bearing capital.

Constitutional Separation of Money and Project Credit

No par guarantee, redemption promise, liquidity backstop, bailout, or emergency support for project notes, capacity bonds, purchase agreements, tranches, or insurance facilities may be transferred—explicitly or implicitly—to the base asset. Losses remain with the contractual credit and equity waterfall. Governance may verify and resolve the claim; it may not monetize it through the bearer asset.

This architecture does not promise cheap capital. It makes the duration holder, default boundary, and loss waterfall legible. If an infrastructure project cannot attract an investor willing to bear its time risk on those terms, the missing object is credit appetite, not money.

Chapter 35

Why a Base Asset Behind These Could Become Money

It is fashionable to say that money is a shared hallucination. That line flatters our cleverness while excusing our passivity. Hallucinations cannot settle debts across adversarial jurisdictions; hallucinations do not finance supply chains. Money works because it is backed by a machine—sometimes a literal machine of war, sometimes a machine of law. This chapter makes the case for a third: verification machines that continuously produce scarce, indispensable utilities.

Post-Bretton-Woods, the machine was compliance. The machine this thesis builds is verification. Cypherpunks did not abolish trust; they automated it. When privacy, proofs, and compute clear across neutral rails, money stops asking for favors and starts paying for facts.

Stated as a conditional, the thesis reads:

> *If* a dense digital civilization continues to rely on AI, global networks, and programmable markets, *and if* states continue to use repression and narrative control rather than explicit default, *and if* the full service path remains stress-deliverable, value accrues without bypass, and a self-custodied loss-bearing holder constituency persists, *then* a bearer base asset **may earn monetary premium** by preserving settlement, proof, compute access, and exit when ordinary substitutes weaken—provided project credit remains separate and without assuming the triad must share one token.

The argument, restated as a conclusion chain:

1. Digital civilization structurally needs privacy, proofs, and compute.
2. These capacities can be standardized and verified through the seven-layer stack.
3. Standardized work can produce cheap-to-check receipts.
4. Receipts represent capacity only where the full path delivers usable service under stress.
5. Stress-deliverable receipts can support markets.
6. Markets can produce recurring fee flows.
7. Fee flows support a scarce base asset only under non-bypassable accrual rules.
8. A persistent, self-custodied, loss-bearing, regime-responsive constituency must warehouse

that asset's risk.

9. Only then, if the asset remains liquid, neutral, verifiable, legally holdable, and agency-preserving, may it earn a store-of-value premium.

The asset may first behave like infrastructure equity—a claim on fee flows from useful services. Only later, if liquidity deepens, neutrality is sustained, and holdability matures across jurisdictions, does it potentially cross the threshold from "productive asset" to "monetary collateral." That transition is not guaranteed; it is the thesis's central conditional claim, and it has a base-rate problem the thesis should state precisely rather than bury or overstate. Stated precisely: **no equity-like claim on a fee stream is known to have made this crossing.** Warehouse receipts, bills of lading, standardized commodity claims, and the utility tokens of Section 34.1.3 — the closest historical analogs to Work Credits and fee units — acquired price and liquidity without ever acquiring monetary premium. The restriction to fee-stream claims matters, because the broader statement — that no asset in monetary history has crossed from use-value to money — is false, and false against the account of commodity money this thesis in fact relies on. Menger [Menger 1892] argued that money arises without design when the most *saleable* good in an economy comes to be accepted for its saleability rather than its use; Mises [Mises 1912] closed the regress by requiring that a money's purchasing power trace back to a prior exchange value it had as a commodity. Silver, salt, grain, colonial Virginia's tobacco, and privately issued banknotes made that crossing. The path they took — use-value, then saleability, then general acceptance — is the path this thesis proposes for a base asset whose use-value is access to verified service, and Menger's account is an argument *for* that path's possibility, not against it. What the record does not contain is a case in which the object making the crossing was a claim on an operator's future fees. Every such claim stayed a claim. The thesis is therefore not describing a developmental stage other assets have passed through, and it is not describing something without precedent either: it is proposing that a bearer asset can take the Mengerian path while a fee stream is routed through it, which is the combination for which the base rate is empty. What follows is the argument for that possibility. Section 30.6.1 supplies the piece Menger requires and this chapter cannot supply on its own — why the saleability bid should attach to the base asset rather than to the typed service claims it prices.

The Mengerian path has a machine-speed variant worth one paragraph. Saleability, in Menger's account, is the property of being accepted without haggling over what the thing is. For a software agent transacting across a trust perimeter that property is not a convenience but a precondition: it cannot investigate a counterparty's unit, so it can accept only units anyone can check (Section 7.5). Acceptance-without-inspection by agents is therefore a candidate first form of the saleability step for an asset whose unit quality is publicly verifiable. This is an argument that the path is open to that constituency, not evidence that anyone has taken it. It adds no precedent: no agent economy has made anything money, the base rate above is unchanged, and Red Line 17 read on agent borrower panels is where the claim would show

up or fail to.

The media-transition literature offers a reading of the fee-stream failure record, and it is offered as a reading, not as evidence: a printed scroll is still a scroll. Each of those precedent systems delivered the *content* of a new capacity—storage, rendering, compute—while keeping the *form* of the old one: a service purchased from an operator, priced as a service, held as an operating claim. Content without structure is exactly what a medium looks like in its first generation [McLuhan 1964], and the utility-token record is that first generation's ledger. On this reading the crossing the thesis proposes is not an increment on those systems but a change of kind—fees are the content; the holder-side stress-deliverable service is the medium—which is why the base rate is empty and why the empty base rate is not dispositive. It is also why the reading earns nothing: no link of Section 8.1 strengthens because a pattern from another century rhymes. The chain still has to prove itself, link by link, in telemetry. The honest summary of the two moves is this. The empty base rate is *evidence against*: it raises the prior that the crossing fails, as the five named precedents and the empty commodity-claim record should. The change-of-kind argument is *the reason the evidence is not dispositive*: it identifies a property the precedents lacked, which changes what the base rate is a base rate *of*. Neither move makes the crossing likely; together they make it testable, which is what the red lines are for.

35.1 As a Store of Value

From the SoV lens, a credible store of value must be credibly scarce, cheap to verify in public, resist censorship and capture, have native demand that is not purely narrative, and avoid being a duration instrument whose real return can be pinned negative by policy.

Privacy, Proofs, and Compute may support a base asset that meets that brief, but the services themselves and their Work Credits are not presumed monetary:

- **Privacy** is purchased because some people and institutions must pay without broadcasting their graph: dissidents, NGOs, treasuries under capital controls, enterprises with sensitive payroll and vendor relationships. In a repression-heavy world, privacy is not a luxury good; it is the hull that keeps savings from becoming an option owned by someone else.
- **Proofs** are purchased because "seeing is believing" has failed. Deepfakes, platform curation, and the liar's dividend make any unproven artifact suspect. Regulated AI and finance regimes require auditable provenance and computation. In that world, proofs are not a niche; they are the affidavit layer of the digital order.
- **Compute** is purchased because intelligence is now a first-class input to production. FLOPs for training, inference, and proving are line items in budgets. Once those FLOPs are wrapped in proofs and standardized as canonical workloads, verified compute becomes a commodity that can be priced, hedged, and stored.

In a world that may choose **stealth default**—negative real yields aided by regulation—over explicit default, a durable SoV must be duration-neutral, peg-proof, stress-deliverable, and held by actors able to bear loss. Recurring utility can support value accrual; it cannot supply moneyness by itself.

The three services above are the use-value. The monetary question is why a holder who wants them under stress would hold the base asset rather than the typed Work Credits that promise them directly, alongside Bitcoin for exit. That question is answered once, at Section 30.6.1, and not restated here: Work Credits are claims on specific capacity that fails in the states the hedge is for, Bitcoin hedges exit and not access, and the base asset is the only workload-agnostic bearer instrument the protocol pays in. The same passage concedes the corollary — that if holders in fact prefer Bitcoin plus Work Credits, or Bitcoin converted on the day, the access experiment and Red Lines 16 and 18 record it, and the thesis treats that as its default null.

Duration-neutrality of the *claim* is not, however, a theory of how plants get built. Section 3.3 and Section 34.5.1 state the completeness condition the rest of this chapter does not repeat: reserve collateral and duration finance are different functions; proofs can audit the second; they must not become it. A stack that supplies Privacy, Proofs, and Compute still requires some institution, outside Layers 0–6, capable of warehousing the construction interval. That institution is not the token.

35.2 As Stack

Seen from the stack angle, the triad is backed not by a metaphor but by a supply chain:

- Layer 0 keeps **machines** honest and powered (verifiable hardware, FERs).
- Layer 1 keeps **packets** flowing under censorship (VerifyReach).
- Layer 2 keeps **code** moving even when app stores and CDNs are hostile.
- Layer 3 keeps **identity** accountable without doxxing.
- Layers 4–5 make **work** and **value** flow through proofs and privacy rails.
- Layer 6 keeps **governance** and **telemetry** legible.

Triad instruments are claims with different seniority and risk. Work Credits represent typed future access or service capacity; FCR, FER, and PIDL artifacts provide evidence; project notes bear duration; LP and staking positions are derivatives or operating claims. Only the base asset is tested as a monetary candidate.

35.3 As Telemetry

Seen from the telemetry angle, every link of the chain has an instrument, and each instrument asks a different question. VERIFYPRICE asks whether anyone can still check the work. Veri-

fyReach and VERIFYSETTLE ask whether checked work can still be reached and paid for under pressure. Delivered Verified Capacity asks whether the service survives the narrowest edge of its physical path (Section 17.9). VerifyFlow asks whether the price is telling the protocol's story or the wrapper's (Section 26.8). Beneath them sit the holder instruments, the Native Monetary Buyer Map and the Principal Map, which ask what none of the others can: whether anyone holds the asset through stress, and whether a human still stands behind the holding (Section 26.8.8, Section 26.8.9).

Together they are the constitution. They keep "trustlessness" from decaying into "trust the custodians," and they keep "store of value" from decaying into "the price went up."

A system that cannot show its own health cannot be money; it can only be marketing.

35.4 Trading One Base Reality for Another

What changes and what stays the same?

What changes: The base reality that backs money. Instead of "the sovereign will repay" or "gold is scarce because geology," the base reality becomes "these capacities are scarce, necessary, and verifiable."

What stays the same: Money is still a claim on work. The work just becomes specific and measurable: proofs that anyone can check, privacy that anyone can use, compute that anyone can verify.

The operative word is **agency**. Monetary arrangements can be read as answers to the question: who has agency over their own savings, payments, and economic identity? Gold answers it with physical possession. Fiat answers it by delegation to institutions. Bitcoin answers it with self-custody and cryptographic keys, and answers it well enough that this thesis builds its settlement layer on top of the answer rather than around it. The triad addresses adjacent surfaces that self-custody does not reach: agency over verification (proofs), over private settlement (paying without being watched), and over productive capacity (compute without permission). Different surfaces, not a later stage (Section 34.1.1). A store-of-value premium, where it exists, accrues to assets that preserve agency on whichever surface is under attack.

Chapter 36

Conclusion: A Bell Labs for Privacy, Proofs, and Compute

The original Bell Labs turned information theory into cables, switches, and semiconductors that quietly reshaped the world. Nobody needed to know the math; they just made phone calls.

The institution this thesis proposes is **Kardashev Labs** (Chapter 32): a build-and-measure lab whose mandate is similar, for a different substrate:

> **Turn privacy primitives, zero-knowledge, and verifiable compute into everyday infrastructure—receipts, rails, and verifiable machines—and test whether a separate bearer base asset can earn monetary premium by preserving agency through that stack.**

That mandate breaks into three concrete programs, and their order matters. The stack and telemetry programs are fundable on their own terms: they produce public infrastructure with buyers who are not token holders. The monetary program is an experiment run on top of them, and the thesis has been explicit that it may fail; nothing in the first two programs is wasted if it does. That is the sense in which the proposal is *stack first, coin optional* (Section 32.1).

1. **Stack program.** Build and harden Layers 0–6: verifiable machines, resilient comms, pseudonymous identity, PoUW and proof factories, privacy corridors, governance and telemetry.
2. **Telemetry program.** Keep VERIFYPRICE, VerifyReach, VERIFYSETTLE, VerifyFlow, Delivered Verified Capacity, the buyer and principal maps, and decentralization metrics live, public, and blunt. Treat dashboards as constitution.
3. **Monetary program.** Test the base asset against the nine-link chain and Native Monetary Buyer Map. Design Work Credits separately as typed, DVC-bounded service or capacity claims; do not ask them to warehouse monetary or project-duration risk.

36.1 A Quiet Arrival

If it does work, the moment of success will not look like an ICO or a conquest. The arrival is quiet.

There will just be a decade in which:

- auditors ask for **receipts**, not screenshots;
- regulators accept **cryptographic attestations** as first-class evidence;
- enterprises budget for **proofs and privacy** the way they once budgeted for bandwidth;
- wallets and apps call `verify(receipt)` as casually as they call `https://`;
- and allocators treat triad assets not as exotic bets, but as part of the boring hedge against repression.

At that point, the thesis will have stopped being a thesis. It will have collapsed into plumbing.

36.2 Two Loops, Not One

One discipline belongs beside that picture, because it is the discipline most likely to be abandoned in a bull market.

The Create/Compute → Prove → Settle → Verify loop tells us whether the cypherpunk monetary stack works. A second loop — Narrative → Wrapper → Allocate/Lever → Hedge → Price → Narrative — tells us how financial markets represent it.

The first loop turns indispensable digital capacity into native monetary value. The second can recognize that value, ignore it, amplify it, or counterfeit its appearance. A serious store-of-value thesis must instrument both. Without VERIFYPRICE, VERIFYREACH, and VERIFYSETTLE, we cannot know whether the stack is real. Without VerifyFlow (Section 26.8), we cannot know whether market price reflects native adoption or external mechanical demand.

Gold can be wrapped without ceasing to be gold. Bitcoin can be held in an ETF without ceasing to settle. But the more monetary use migrates from bearer ownership into custodial and synthetic claims, the more the financial representation can diverge from the very properties that justified the premium. The destination is not merely a higher price. It is a monetary object whose price, utility, custody, and settlement remain connected closely enough that each can still discipline the others.

> VERIFYPRICE, VERIFYREACH, and VERIFYSETTLE verify the stack. VerifyFlow verifies the story the market is telling about it. Only the holders, by how they hold, can verify the money.

36.3 The Question Under Test

Bitcoin established that **digital scarcity alone** can be money. That result stands, this thesis does not revise it, and nothing here improves on the mechanism that produced it (Section 34.1.1). The question taken up in these pages is a different one: whether **native utility that the world must keep buying**, secured by proofs anyone can check, can also carry monetary premium—or whether it only ever prices as a service.

Privacy must preserve agency.

Proofs must bound what can be faked—not create truth, but shrink the territory where forgery pays.

Compute must power intelligence—verifiable and honest, not merely performed.

When these three clear across neutral settlement, money stops being an article of faith and becomes what it always wanted to be: **a record of work that cannot be faked and does not need permission to move**.

In the decade ahead we will learn whether digital necessities, wrapped in proofs anyone can verify, can become candidate money for a world where verification is more valuable than permission.

36.4 Equip, Don't Manage

The purpose of the stack is not to make people more legible to institutions. It is to preserve the capacity to act when institutions become coercive, synthetic, or brittle. A CBDC can manage citizens. A custodial stablecoin can manage customers. A platform AI can manage users. A cypherpunk monetary stack must do the opposite: **equip actors**.

Privacy equips people to move value without becoming visible by default. Proofs equip people to verify claims without trusting platforms. Compute equips people and agents to purchase machine work without asking permission from hyperscalers. Telemetry equips the public to see drift before the stack becomes another soft guarantee.

Gold condenses geology. Bitcoin condenses thermodynamics. The triad condenses **verifiable agency**: the ability to pay, prove, and compute under adversarial conditions. Three hard facts, held in the present tense, none of them retiring another—and the third is the only one still under test.

This closes the circle opened in Section 3.6. Fiscal dominance makes financial repression attractive; the marginal buyers of the reserve asset have become rules and the leverage that serves them; AI makes truth and cognition platform-mediated; household non-participation makes administrative management politically tempting; and physical bottlenecks make compute, power, and hardware strategically scarce. Gold, silver, Bitcoin, and hard assets remain the bridge (Sections 29.1 and 29.1.1). Privacy, Proofs, and Compute can support an associ-

ated base-asset monetary candidate only when the full nine-link chain proves, instrument by instrument, that the service stays deliverable when it is attacked, that demand cannot route around the asset without paying for it, and that the people holding it can bear the loss. Every instrument named in this thesis exists to make one of those three sentences measurable.

And the circle has a household edge the sovereign layers do not supply. Proof tells us whether a claim is true; representation tells us whom the claim describes; capability tells us what participants can actually do as a result. A stack that verifies everything and represents almost no one is a more efficient ledger for the same exhaustion dynamics it was built to escape, which is why the representation discipline travels with every aggregate this thesis publishes and the capability layer travels with the ninth requirement. Equipping actors, in the end, means more than non-custodial keys: it means the architecture does not mistake extraction for resilience because it measured itself in dollars rather than people.

It also means that the actor stays a person. As more of what this stack sells is bought by agents acting several hops from anyone, ownership moves from the balance to the key root: whether a human can still revoke, cap, and re-key the processes that act in their name (Section 7.6). A stack that verifies every agent and answers to no one has equipped the agents and managed the people.

The thesis fails if the dashboards fail, if verification centralizes, if privacy rails cannot stay lawful and usable, or if value does not accrue to the asset. Those are not escape clauses; they are the terms of the bet.

> The question this stack answers is not whether its ledgers balance. It is whether the people inside them are owners or tenants.

Sources

Each entry opens with a bracketed key, [Author Year], and the body cites by the same key: either the key itself where a work, dataset, institution, or case is named, or a linked year where the prose already reads "Author (Year)." Both forms are hyperlinks to the entry. Keys carry letter suffixes where one author has several items in a year, in the order listed here. Entries are grouped by the part of the argument they serve, not alphabetically. Entries marked *landscape reference* document the state of the field the thesis is positioned against and carry no body key; every other entry is cited at least once.

Theory & Framing

[Davidson & Rees-Mogg 1997] **Davidson, J. D., & Rees-Mogg, W.** (1997). *The Sovereign Individual: How to Survive and Thrive During the Collapse of the Welfare State*. New York: Simon & Schuster / Touchstone. https://www.simonandschuster.ca/books/The-Sovereign-Individual/James-Dale-Davidson/9780684832722. *Cited in Chapter 1 for the prediction that networked computation erodes centralized control and provokes capital controls and nationalization. Framing only; no quantity or forecast is adopted.*

[Booth 2020] **Booth, J.** (2020). *The Price of Tomorrow: Why Deflation is the Key to an Abundant Future*. Stanley Press. https://thepriceoftomorrow.com/. *Cited in Section 11.1 for technological deflation as a secular force on unit compute cost. Framing, not measurement.*

[McLuhan 1964] **McLuhan, M.** (1964). *Understanding Media: The Extensions of Man*. New York: McGraw-Hill. https://archive.org/details/understandingmed0000mclu. *Cited for framing only: that a medium's structural effects arrive as externalities of adoption rather than as features anyone selects, and that the content of a new medium is initially the old medium's content. Used in Section 2.0.1 and the Executive Memo as pattern, not as evidence; no claim of this thesis depends on it, and nothing in it is offered as measurement.*

[McLuhan 1962] **McLuhan, M.** (1962). *The Gutenberg Galaxy: The Making of Typographic Man*. Toronto: University of Toronto Press. https://archive.org/details/gutenberggalaxym0000mclu. *Cited for the sensorium framing—that a medium restructures the ratio of the senses before anyone notices—and for print's restructuring of the private reader. Same limits as above: synthesis, not experiment; mechanism, not quantities.*

[Eisenstein 1979] **Eisenstein, E. L.** (1979). *The Printing Press as an Agent of Change: Communica-*

tions and Cultural Transformations in Early-Modern Europe. Cambridge: Cambridge University Press. https://www.cambridge.org/core/books/printing-press-as-an-agent-of-change/7DC19878AB937940DE13075FE839BDBA. *Cited in Section 2.0.1 for the standardization thesis: that typographic uniformity preceded and enabled the standardization of physical production, and that print's consequences were causal and profound. It is cited for mechanism only. The thesis does not adopt any necessity claim—that the industrial revolution was possible only because of movable type—and no argument here rests on one.*

[Ong 1982] **Ong, W. J.** (1982). *Orality and Literacy: The Technologizing of the Word*. London: Methuen. https://archive.org/details/oralityliteracyt00ongw. *Cited for interiority: that sustained silent reading is a technology-dependent practice and the basis of the private self as an epistemic unit. Used in Chapter 9 to place private settlement in a longer arc. Framing only; the thesis's claim about payments as the load-bearing privacy surface is unchanged by it.*

[Anderson 1983] **Anderson, B.** (1983). *Imagined Communities: Reflections on the Origin and Spread of Nationalism*. London: Verso. https://archive.org/details/imaginedcommunit0000ande_s3l3. *Cited for one connection only: that print created mass simultaneous identity, and that the Native Monetary Buyer Map is how a monetary constituency renders itself legible to itself. The analogy is structural; no claim about adoption is derived from it.*

[Culkin 1967] **Culkin, J. M., S.J.** (1967, March 18). A Schoolman's Guide to Marshall McLuhan. *The Saturday Review*, pp. 51–53, 71–72. http://www.unz.org/Pub/SaturdayRev-1967mar18-00051. Open copy: https://webspace.royalroads.ca/llefevre/wp-content/uploads/sites/258/2017/08/A-Schoolmans-Guide-to-Marshall-McLuhan-1.pdf. *Cited for attribution: "We become what we behold. We shape our tools and thereafter they shape us" is Culkin's line, routinely misattributed to McLuhan. It appears in this thesis, if at all, as Culkin's.*

[Perez 2002] **Perez, C.** (2002). *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*. Cheltenham: Edward Elgar. https://doi.org/10.4337/9781781005323. *Cited for one mechanism only: that the installation phase of an infrastructure buildout is financed by structures that outrun the productive economy's ability to absorb them, and that the resulting clearing event reprices the installed base — transferring it from weakly financed owners to stronger users at lower capital cost — which is what makes the deployment phase's cheap substrate. Used in Section 29.2 (the capital cycle as a cross-cutting clock on the four states), Section 29.1 (the window closing by event rather than accrual), Section 13.3.2 and Section 11.1 (financial repricing as a deflation channel distinct from technological deflation), Section 29.3 (the fee base as a deployment-phase quantity), and Section 33.9 (the rescue-born pathway to State 4). The thesis does not adopt the historical periodization (irruption/Frenzy/Turning Point/Synergy/Maturity as dated phases), does not claim AI constitutes a "sixth revolution" or adopt Perez's own current reading of AI as a third boost to the ICT surge, does not use the framework to forecast any asset, sector, or macro variable, and no red line, threshold, or falsification condition depends on it. Where the frame and the thesis's instruments*

conflict, the instruments govern.

[Gorton & Pennacchi 1990] **Gorton, G., & Pennacchi, G.** (1990). Financial Intermediaries and Liquidity Creation. *The Journal of Finance*, 45(1), 49–71. https://doi.org/10.1111/j.1540-6261.1990.tb05080.x. *Cited in Section 13.3.10 for one mechanism: that a security can be designed so that no counterparty gains from producing private information about it, and that such a security circulates at par for that reason. Mechanism only; no magnitude and no claim about bank liabilities is imported, and no result of this thesis depends on the model.*

[Dang et al. 2020] **Dang, T. V., Gorton, G., & Holmström, B.** (2020). The Information View of Financial Crises. *Annual Review of Financial Economics*, 12, 39–65. https://doi.org/10.1146/annurev-financial-012820-025249. *Cited in preference to the widely circulated but undated "Ignorance, Debt and Financial Crises" working-paper lineage from which it derives, because a citation record should not rest on a moving target. Cited for the information-insensitivity account of debt: that the money-like security is the one whose payoff is least sensitive to private information, and that a run is the transition from information-insensitive to information-sensitive rather than a revision of fundamentals. Mechanism only. The thesis explicitly does* not *adopt the corollary that opacity is required — Section 13.3.10 argues an opposite route to the same property — and that departure is this thesis's claim, not theirs, and carries its own falsifier.*

[Holmström 2015] **Holmström, B.** (2015). Understanding the Role of Debt in the Financial System. BIS Working Paper No. 479. https://www.bis.org/publ/work479.htm. *Cited for the "no questions asked" formulation and for the position that opacity is a design feature of money markets rather than a defect. Used as the statement of the view Section 13.3.10 inverts; the inversion is ours and is falsifiable through Red Line 17.*

[Gorton & Metrick 2012] **Gorton, G., & Metrick, A.** (2012). Securitized Banking and the Run on Repo. *Journal of Financial Economics*, 104(3), 425–451. https://doi.org/10.1016/j.jfineco.2011.03.016. *Cited for one empirical characterization: that the 2007–2009 funding collapse ran through collateral quality and haircuts rather than through a revision of fundamentals participants did not already hold. Used to identify* which *informational margin failed. One descriptive range from it — structured-collateral repo haircuts rising from near zero to 40–45% — is quoted for illustration in Red Line 17; no red-line threshold depends on it, and no other quantity from it enters this thesis, and the analogy is to the margin, not to the setting.*

[Menger 1892] **Menger, K.** (1892). On the Origin of Money. *The Economic Journal*, 2(6), 239–255. Translated by C. A. Foley. https://doi.org/10.2307/2956146. Free full text: https://mises.org/library/book/origins-money. *Cited in Chapter 35 for the saleability account of commodity money: that money arises without design when the most saleable good comes to be accepted for its saleability rather than its use, and that the path runs use-value → saleability → general acceptance. Used to correct an earlier overstatement that no asset had crossed from use-value to money, and to state that the empty base rate is for equity-like claims on a fee stream specifically. Menger's account is treated*

as an argument for the possibility of the path this thesis proposes, not as evidence that any particular asset will take it.

[Mises 1912] **Mises, L. von** (1912/1953). *The Theory of Money and Credit*. Translated by H. E. Batson. New Haven: Yale University Press (original German edition 1912). https://mises.org/library/book/theory-money-and-credit. *Cited in Chapter 35 for the regression theorem only: that a money's present purchasing power traces back to a prior exchange value it held as a commodity. Used alongside Menger (1892) to fix the scope of the base-rate claim. No quantity, forecast, or policy position from the work is adopted, and the thesis does not rely on its broader monetary theory.*

[Doepke & Schneider 2017] **Doepke, M., & Schneider, M.** (2017). Money as a Unit of Account. *Econometrica*, 85(5), 1537–1574. https://doi.org/10.3982/ECTA13541. *Cited in Section 13.7.3 for one mechanism: that the choice of unit of account is a coordination problem embedded in contract networks, and therefore exhibits inertia that portfolio choices do not. Used for stickiness and strategic complementarity; no result, magnitude, or welfare claim is imported, and the thesis's denomination instrument is not derived from the model.*

[Grossman & Stiglitz 1980] **Grossman, S. J., & Stiglitz, J. E.** (1980). On the Impossibility of Informationally Efficient Markets. *The American Economic Review*, 70(3), 393–408. https://www.jstor.org/stable/1805228. *Cited in Section 13.3.10 for one structural result: that costly information cannot be fully reflected in prices, because an agent who earns no return on producing it will not produce it. Used to state an objection to this thesis's own mechanism, and to identify what answers it — that a proof system separates the party bearing the cost from the parties consuming the result, so the expenditure is voluntary rather than rent-financed. No model, magnitude, or efficiency claim is imported; this thesis takes no position on market efficiency, and the noise-trader apparatus of the paper plays no role here.*

[Baumol 1982] **Baumol, W. J.** (1982). Contestable Markets: An Uprising in the Theory of Industry Structure. *The American Economic Review*, 72(1), 1–15. https://www.jstor.org/stable/1831298. *Cited in Section 13.3.10 for one mechanism only: that the credible possibility of entry can discipline conduct without any entry being observed. Used as the structural analogue for verification that is continuously available and almost never exercised. No result about industry structure, sunk costs, contestability's empirical standing, or welfare is imported, and no threshold in this thesis is derived from it.*

[Kaldor 1939] **Kaldor, N.** (1939). Speculation and Economic Stability. *The Review of Economic Studies*, 7(1), 1–27. https://doi.org/10.2307/2967593. *Cited in Section 13.3.11 for one general mechanism only: that a service flow accruing to the holder of a physical good, rather than to the holder of a claim on it, can carry a price a cash-flow model does not reproduce. It is not cited for gold, which is a near-zero-convenience-yield commodity whose price residual is conventionally attributed to monetary and portfolio demand rather than to storage convenience; and it is not a claim that a triad asset is a storage commodity. No magnitude in this thesis derives from it.*

[Working 1949] **Working, H.** (1949). The Theory of Price of Storage. *The American Economic Review*, 39(6), 1254–1262. https://www.jstor.org/stable/1816601. *Companion to Kaldor, cited for the same mechanism and subject to the same limits: the storage-theory account of convenience yield in consumption commodities, where the service flow is an option against stockout. The analogy in Section 13.3.11 is to the mechanism, not to the setting.*

Market Structure, Flows & Inelastic Markets

[Green 2026a] **Green, M. W.** (2026, July 26). A Semi-Theory of Almost Everything: Who Was Actually Setting Semiconductor Prices This Spring? Tier1 Alpha. Paywalled. https://www.yesigiveafig.com/p/a-semi-theory-of-almost-everything. *The flow-elasticity, holder-recycling, and leveraged-rebalancing machinery in Appendix H is adapted from this work; the generalization to monetary telemetry is ours, as are any errors in it. The daily-reset rebalancing identity and the volatility-drag result that the appendix builds on are older than this essay and are cited to Cheng & Madhavan (2009) and Avellaneda & Zhang (2010) below.*

[Nakamoto 2008] **Nakamoto, S.** (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. https://bitcoin.org/bitcoin.pdf. *Cited for the proof-of-work construction that mints digital scarcity from hash puzzles and for the seventeen-year record the thesis uses as the control case throughout.*

[Shannon 1948] **Shannon, C. E.** (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3), 379–423; 27(4), 623–656. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x. *Cited once, in the introduction's Bell Labs framing: the canonical case of an institution turning theory into deployed infrastructure.*

[Green 2026b] **Green, M. W.** (2026, August 16). *Anchors Aweigh, My Boys.* "Yes, I give a fig…" (Substack). Paywalled. https://www.yesigiveafig.com/p/anchors-aweigh-my-boys. *The automatic-versus-marginal buyer distinction, the observation that market-value-weighted bond indexes absorb less DV01 per dollar after a selloff, and the claim that the quality of long-duration demand has deteriorated are drawn from this work and used in Section 3.1.3, Section 3.3, Section 26.8.7, and Section 26.8.8. Commentary rather than peer-reviewed research; used for mechanism, not for quantities. Tactical claims and the attribution of any particular yield move uniquely to duration withdrawal are listed and excluded in Appendix I.4.*

[Green 2026c] **Green, M. W.** (2026, May 24). *What the Treasury Needs.* "Yes, I give a fig…" (Substack). Paywalled. https://www.yesigiveafig.com/p/what-the-treasury-needs. *Used selectively for the distinction among market repair, fiscal management, financial repression, and administrative enclosure, and for the proposition that policy can change duration-bearing capacity without erasing economic loss. Commentary rather than peer-reviewed research; used for mechanism, not for the proposal's illustrative quantities or for a claim that any facility is necessary.*

[Green 2026d] **Green, M. W.** (2026, August 19). *Sometimes, You Get What You Need.* "Yes, I give a fig…" (Substack). Paywalled. https://www.yesigiveafig.com/p/sometimes-you-get-what-you-need. *Used as a scenario-framing example for the market-repair branch in Section 29.2. Its tactical sequence for rates, currencies, gold, mortgages, and equities is quarantined rather than adopted as forecast.*

[Green 2026e] **Green, M. W.** (2026, August 30). *The Vibecession Was Real.* "Yes, I give a fig…" (Substack). Paywalled. https://www.yesigiveafig.com/p/the-vibecession-was-real. *Used for mechanism, not quantities: the dollar-weighted versus person-weighted aggregation distinction and its covariance identity (adopted as the representation discipline on board aggregates, Section 26.8.8); the price-level-versus-inflation-process-versus-balance-sheet-hysteresis decomposition and the consumption-through-depletion mechanism (adopted in the household-transmission observables of Section 29.2 and the exhaustion gradient of Section 3.6); the state-versus-household resilience distinction $R_{state} \neq R_{household}$ and the formal operating-surplus/resilience-stock model (adopted in Section 33.9); and the lived-solvency framing behind the incumbency objection (Section 34.3.6). Commentary rather than peer-reviewed research. Excluded: the essay's illustrative simulation percentages (presented there as mathematical possibility, not prevalence), its pricing-power magnitudes (anecdotal), any psychological attribution to named officials, and its secondary statistics, which are treated as provisional pending tracing to underlying series.*

[Green 2026f] **Green, M. W.** (2026). *The Quiet Turning.* "Yes, I give a fig…" (Substack). Paywalled. https://www.yesigiveafig.com/p/the-quiet-turning. *Used for one observation and two objections, each of which reads against this thesis rather than for it.* ***The observation****: that cheap checking routinely goes unperformed — the record public, the analysis absent — which refuted the original "everyone has already looked" formulation of Section 13.3.10 and forced its restatement as deterrence on the claimant rather than audit by the counterparty.* ***The objections****: that the free-rider structure of costly information applies to honesty as it does to prices, and that discovery without consequence changes nothing. The first is answered in Section 13.3.10 from the separation between the party bearing the cost of proof and the parties consuming it; the second is instrumented as the plain-language reading of Red Line 14 (Chapter 30).* ***What is not drawn from this work****: the contestability framing that replaces the retracted formulation is cited to Baumol (1982) and the free-rider structure to Grossman & Stiglitz (1980); the monetary consequence, the adversarial-invariance claim, the scope bound on computable claims, and the synthesis of any of it with the collateral literature are this thesis's own and are endorsed by no one. Commentary rather than peer-reviewed research, and published behind a paywall — which is why no argument here rests on it: every mechanism it prompted is carried by an independently citable source a reader can open. Excluded: its fraud aggregates, which are self-reported, overlap across FTC and IC3 collection, and are not summable, a limitation the essay states itself; its generational characterization, which is anecdotal; and its attributions of motive to named individuals and firms.*

[Green 2026g] **Green, M. W.** (2026, September 20). *The Market Has a Rule of Its Own: Before the Fed Reads the Long End as a Message, It Should Identify Who Is Speaking.* "Yes, I give a fig…" (Substack). Paywalled. https://www.yesigiveafig.com/p/the-market-has-a-rule-of-its-own. *Used in Section 3.1.3 for one formulation—description becomes instruction: a market-value-weighted index reports the market, and once capital is managed against it the report becomes the order—with its corollary that the seller of credit comes to determine the portfolio of the buyer. The decomposition Section 13.7 describes—net demand for Treasury interest-rate risk by tenor, in DV01, by buyer and decision rule, with a published residual—is the one the essay asks the Board to run. The mechanism itself is cited to the two staff papers the essay rests on ([Barth et al. 2024], [Kruttli et al. 2021]) rather than to the essay. Commentary; mechanism, not quantities. No figure from the essay is used. Excluded and listed in Appendix I.4: the closing claim that mandate-followers and their arbitrageurs are the market's dominant voice, which the essay's own proposed decomposition has not yet tested; and any attribution of a policy move or a yield level to the mechanism. Seven essays by one commentator are drawn on in this document, and the reader should weigh that. The rule applied to all seven: an essay is cited where the thesis first adopts one of its constructs, in the quarantine table where its claims are held back, and in Appendix J where it prompted a revision; everywhere else the text points to its own section. No essay is used for a quantity or a forecast.*

[Gabaix & Koijen 2021] **Gabaix, X., & Koijen, R. S. J.** (2021). In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis. NBER Working Paper No. 28967. https://www.nber.org/papers/w28967. *Cited for the inelastic-markets premise that flows move prices at all, not for the market-impact function used in Appendix H; that square-root form comes from the market-microstructure execution-impact literature.*

[Cheng & Madhavan 2009] **Cheng, M., & Madhavan, A.** (2009). The Dynamics of Leveraged and Inverse Exchange-Traded Funds. *Journal of Investment Management*, 7(4), 43–62. https://joim.com/article/the-dynamics-of-leveraged-and-inverse-exchange-traded-funds/. Preprint: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1393995. *Cited in Appendix H for the daily-reset rebalancing identity — that a fund with leverage L must trade $L(L-1)$ times its assets times the underlying return to restore target exposure, and that this holds for inverse as well as long leverage — and for the observation that both long and inverse leveraged funds therefore trade in the direction of the underlying move. The recycling boundary and the Wrapper Recycling Ratio built on that identity are ours.*

[Avellaneda & Zhang 2010] **Avellaneda, M., & Zhang, S.** (2010). Path-Dependence of Leveraged ETF Returns. *SIAM Journal on Financial Mathematics*, 1(1), 586–603. https://doi.org/10.1137/090760805. *Cited in Appendix H.9 for the continuous-time expression of a daily-reset leveraged fund's compound return in terms of the underlying's return and realized variance — the source of the $-L^2\sigma^2/2$ variance penalty and of the statement that the terminal value is path-dependent. The financing and expense terms added there, and the classification consequence drawn in Table 8.4, are*

ours.

[Messari 2024–26] **Messari.** (2024–2026). Protocol revenue and token issuance series for Filecoin, Render, Akash, Golem, and Livepeer. Messari analytics dashboards. https://messari.io/project/filecoin. *Cited in Section 34.1.3 as one of the two public sources for the observation that native fee revenue on the closest utility-token precedents has covered low single digits of issuance. Dashboard series, not audited financials; the ratio depends on each dashboard's classification of fees, burns, and incentive emissions, and the thesis relies on the order of magnitude only.*

[Chainalysis 2024] **Chainalysis.** (2024, October). *The 2024 Geography of Cryptocurrency Report*. https://www.chainalysis.com/wp-content/uploads/2024/10/the-2024-geography-of-crypto-report-release.pdf. *Cited in Section 13.5.1 for one qualitative finding: that in the emerging-market currency episodes it documents (Argentina, Turkey, Nigeria), dollar stablecoins rather than bitcoin or privacy assets were the instrument of savings flight and cross-border settlement. Industry analytics, not audited data; used for the direction of revealed preference, not for any volume figure.*

[Chainalysis 2020] **Chainalysis.** (2020, June). 60% of Bitcoin Is Held Long Term as Digital Gold. What About the Rest? Chainalysis blog. https://blog.chainalysis.com/reports/bitcoin-market-data-exchanges-trading/. *Cited in Section 12.6.1 for the estimate that 2.3–3.7 million bitcoin have not moved in five years or more and are treated as lost. Used to anchor a modelling assumption by loose analogy only; the thesis's 1%-per-year figure is not derived from it.*

[Token Terminal 2024–26] **Token Terminal.** (2024–2026). Fees, revenue, and token incentives series for Filecoin, Render, Akash, Golem, and Livepeer. Token Terminal analytics dashboards. https://tokenterminal.com/explorer/markets/depin. *Cited in Section 34.1.3 alongside Messari for the same observation and under the same caveat. Where the two dashboards disagree, the thesis does not adjudicate; it states the order of magnitude both support.*

Energy, Industrial Capacity & Sovereign Optionality

[Doomberg 2026a] **Doomberg.** (2026, July 29). *Fire Horse*: analysis of Chinese energy strategy, coal, refining, conversion capacity, and stockpiling. Doomberg (Substack). Paywalled. https://newsletter.doomberg.com/p/fire-horse. *Used in Chapter 33 to characterize a strategy of conversion optionality and "resilience over efficiency." This is an analyst presentation, not peer-reviewed work; figures cited are its accounting rather than independent measurement, and no claim in this thesis is load-bearing on them.*

[Doomberg 2026b] **Doomberg.** (2026, August 9). *Endangered Specious.* Doomberg (Substack). Paywalled. https://newsletter.doomberg.com/p/endangered-specious. *Used selectively for the mechanism that gross feedstock is not deliverable output when conversion capacity is the binding cut. Commentary rather than peer-reviewed research; used to motivate the Topological Scarcity Lemma*

and DVC, not for article-specific quantities or a conclusion that aggregate abundance guarantees local resilience.

[Doomberg 2026c] **Doomberg.** (2026, August 13). *A Minor Detail.* Doomberg (Substack). Paywalled. https://newsletter.doomberg.com/p/a-minor-detail. *Used selectively as a stress scenario illustrating common-cause dependence across fuel, logistics, cooling water, and generation. The geopolitical sequence is not treated as a base case or tactical forecast, and no unsourced quantity enters the model.*

[Doomberg 2026d] **Doomberg.** (2026, August 18). *Over Target.* Doomberg (Substack). Paywalled. https://newsletter.doomberg.com/p/over-target. *Used selectively for the principal–agent mechanism in the Sponsor Dependency Board: a subsidized operator and its sponsor may have different terminal objectives. Claims about the intent behind any particular attack or escalation path are not adopted.*

[Gromen 2025–26] **Gromen, L.** *The Forest for the Trees* (FFTT LLC). FFTT research notes, 2025–2026 (paywalled); specific issues cited in text. *Three framings are drawn from this work: sovereign debt as a claim on future real surplus, with monetary credibility jointly determined by productive capacity and the state's ability to mobilize it (Section 33.1); the reflexive dependence of fiscal capacity on collateral values (Section 3.2); and the reshoring trilemma (Section 33.3). Commentary rather than peer-reviewed research; used for framing, not for quantities. Quantitative claims from the same source that could not be verified are listed and excluded in Appendix I.4.* https://fftt-llc.com/.

[NERC annual] **North American Electric Reliability Corporation (NERC).** (annual). *Long-Term Reliability Assessment*. Atlanta, GA: NERC. https://www.nerc.com. *Cited in Section 17.10.3 as the regulatory anchor of the resource-adequacy literature: the standing, published assessment of how much reserve capacity the North American grid carries against demand and how that margin is trending. The thesis imports the framing — redundancy priced against the cost of idle capacity, and resilience that is not metered becoming rent — not any figure, and the FCR/DVC instruments are its own.*

Note on the synthesis. The layered reading in Chapter 33—physical throughput beneath sovereign balance sheet beneath market realization beneath verification—is our construction, assembled from the above together with the market-structure sources in the preceding section. Sovereign Optionality adapts pathway diversity to facility and network scale; Delivered Verified Capacity restores the conversion-throughput term as scenario-specific service max flow. The Topological Scarcity Lemma, Pressure–Capacity Corridor, and their application to AfterFiat are our synthesis. Errors in that adaptation are ours.

[Bessent 2026] **Bessent, S.** (2026, June 23). *American Economic Statecraft in the 21st Century*. Remarks at the Economic Club of New York's America 250 Gala Dinner. U.S. Department of the Treasury, press release sb0539. https://home.treasury.gov/news/press-releases/sb0539.

A primary source, cited in Section 33.3 for stated doctrine only. Quotations are verbatim. A policy speech is evidence of intent and framing, not of outcomes, and this document draws no inference from it about what will actually be built.

[Kimi Team 2026] **Kimi Team.** (2026). *Kimi K3: Open Frontier Intelligence*. arXiv:2607.24653. https://arxiv.org/abs/2607.24653. *Cited in Section 33.4 for a single narrow proposition: that openly released weights have approached the proprietary frontier closely enough to matter economically. The paper reports the model trailing the strongest proprietary systems overall, and that qualification travels with the citation. No claim about inference pricing is drawn from it.*

Macro: Debt, Financial Repression, Real Yields

[Gaspar et al. 2025] **Gaspar, V., Gonçalves, C. E., & Poplawski-Ribeiro, M.** (2025, September 17). Global Debt Remains Above 235% of World GDP. *IMF Blog / Global Debt Monitor*. International Monetary Fund. https://www.imf.org/en/blogs/articles/2025/09/17/global-debt-remains-above-235-of-world-gdp. *Cited in Section 3.1 for the headline global debt-to-GDP multiple only.*

[IMF 2025] **International Monetary Fund.** (2025). *Global Debt Monitor 2025*. IMF Global Debt Database. https://www.imf.org/external/datamapper/datasets/GDD. *Cited in Section 3.1 as the dataset behind the debt figures; no projection beyond those stated in the text is adopted.*

[IMF 2024] **International Monetary Fund.** (2024, October). *Fiscal Monitor, October 2024: Putting a Lid on Public Debt*. IMF. Series page: https://www.imf.org/en/Publications/FM. [Reports global public debt exceeding $100 trillion in 2024, approaching 100% of global GDP by the end of the decade; cited in Part I for the headline stock figure. Superseded for the ratio projections by the April 2026 edition below.] *Cited for the stock figure only.*

[IMF 2026] **International Monetary Fund.** (2026, April). *Fiscal Monitor, April 2026*. IMF. Series page: https://www.imf.org/en/Publications/FM. [Reports global public debt at just under 94% of GDP in 2025, projected to reach 100% by 2029.] *Cited in Section 3.1 for the public-debt ratio and its end-of-decade projection; an official projection, quoted as such.*

[Reinhart & Sbrancia 2015] **Reinhart, C. M., & Sbrancia, M. B.** (2015). The Liquidation of Government Debt. *IMF Working Paper 15/7*. International Monetary Fund. https://www.imf.org/external/pubs/ft/wp/2015/wp1507.pdf. *Source of the "liquidation effect" construction used in Appendix A, Definition 7: the annual real transfer from bondholders to the sovereign when real rates are held negative, and its historical size in advanced economies over 1945–1980.*

[Barth et al. 2025] **Barth, D., Beltran, D., Hoops, J., Kahn, R. J., Liu, Z., & Perozek, M.** (2025, October 15). The Cross-Border Trail of the Treasury Basis Trade. *FEDS Notes*. Board of Governors of the Federal Reserve System. Series: https://www.federalreserve.gov/econres/n

otes/feds-notes/. [The three figures this thesis relies on: Cayman-domiciled hedge funds' Treasury holdings of roughly $1.85 trillion at end-2024; an understatement of hedge-fund Treasury holdings in the TIC data of roughly $1.4 trillion; and a resulting overstatement of the U.S. household saving rate of roughly 2.1 percentage points over 2023–24.] *Cited in Section 3.1.3 for the scale and financing of Cayman-domiciled hedge-fund Treasury positions and for the resulting mismeasurement in the TIC data and the Financial Accounts. Staff research, not a Board policy position. The note itself cautions that the reported gap is not attributable solely to the cash–futures basis trade, and that caution is reproduced in the main text.*

[Barth et al. 2024] **Barth, D., Kahn, R. J., Monin, P., & Sokolinskiy, O.** (2024, June). Reaching for Duration and Leverage in the Treasury Market. *Finance and Economics Discussion Series* 2024-039. Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/econres/feds/reaching-for-duration-and-leverage-in-the-treasury-market.htm. [The findings this thesis relies on: variation in mutual funds' Treasury-futures positions drives much of the time-series variation in aggregate Treasury-futures open interest, including more than sixty percent of the rise examined; those positions are largely associated with funds reaching for benchmark duration while holding cash portfolios tilted toward higher-return, lower-duration assets; and hedge funds frequently take the other side, adding leverage.] *Cited in Section 3.1.3 as the primary source for the chain from benchmark duration-matching through futures to leveraged intermediation. Staff research, not a Board policy position. It does not establish that benchmark mechanics set the level of any yield, and is not used for that.*

[Kruttli et al. 2021] **Kruttli, M. S., Monin, P. J., Petrasek, L., & Watugala, S. W.** (2021, June). Hedge Fund Treasury Trading and Funding Fragility: Evidence from the COVID-19 Crisis. *Finance and Economics Discussion Series* 2021-038. Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/econres/feds/hedge-fund-treasury-trading-and-funding-fragility-evidence-from-the-covid-19-crisis.htm. [The findings relied on: hedge-fund gross Treasury exposures roughly doubled from 2018 to February 2020, to about $2.4 trillion; funds concentrated in the cash–futures basis trade faced greater margin pressure and reduced Treasury exposures and repo borrowing the most during March 2020.] *Cited in Section 3.1.3 for the stressed-regime behavior of leveraged Treasury intermediation. Staff research, not a Board policy position. Not used to attribute the March 2020 episode to benchmark demand, which neither paper claims.*

[Rose 2021] **Rose, J.** (2021). Yield Curve Control in the United States, 1942 to 1951. *Economic Perspectives*, Federal Reserve Bank of Chicago, 45(2). https://www.chicagofed.org/publications/economic-perspectives/2021/2. PDF: https://www.chicagofed.org/-/media/publications/economic-perspectives/2021/ep2021-2-pdf.pdf. *Cited in Section 3.1 and Section 3.1.1 for the 1942–1951 bill peg, the long-bond cap, and the 1951 Accord. Historical description; no inference about future policy is drawn from it.*

Social / Disinformation / Deepfakes / Provenance

[Twitter Files 2022–23] **"Twitter Files."** (2022–2023). Series of internal document releases published by independent journalists via the Twitter/X platform. [Primary documents; this thesis cites the releases as evidence of government–platform coordination, not as dispositive proof of any particular claim.] https://twitterfiles.substack.com/. *Cited in Section 3.5 as evidence that coordination occurred; no individual claim in the releases is treated as established.*

[House Judiciary 2023] **U.S. House Committee on the Judiciary (Republican Staff).** (2023, June 26). *The Weaponization of CISA: How a "Cybersecurity" Agency Colluded with Big Tech and "Disinformation" Partners to Censor Americans*. Interim Staff Report. https://judiciary.house.gov/sites/evo-subsites/republicans-judiciary.house.gov/files/evo-media-document/cisa-staff-report6-26-23.pdf. *Cited in Section 3.5 as a partisan staff report that points to primary documents; evidence of public discourse, not an authority on any finding.*

[State Dept. 2021–25] **U.S. Department of State.** (2021–2025). Global Engagement Center (GEC): About Us. U.S. Department of State. https://2021-2025.state.gov/about-us-global-engagement-center-2/. *Cited in Section 3.5 as the agency's own description of its mandate; primary source for the program's existence, not for its effect.*

[Harwell & Patton 2025] **Harwell, D., & Patton, D.** (2025, October 22). We uploaded a fake video to 8 social apps. Only one told users it wasn't real. *The Washington Post*. [Seven of eight platforms tested stripped or failed to surface the C2PA Content Credentials attached to the upload; cited in Chapter 3.] https://www.washingtonpost.com/technology/2025/10/22/ai-deepfake-sora-platforms-c2pa/. *One platform sweep on one date; cited for the stripping result, not as a base rate for platform behaviour.*

[UNESCO 2025] **UNESCO.** (2025). *Deepfakes and the Crisis of Knowing*. UNESCO briefing. *Cited in Chapter 3 for the "crisis of knowing" framing.* https://www.unesco.org/en/articles/deepfakes-and-crisis-knowing.

[Ahmed et al. 2024] **Ahmed, S., Bee, A. W. T., Ng, S. W. T., & Masood, M.** (2024). Social Media News Use Amplifies the Illusory Truth Effects of Viral Deepfakes: A Cross-National Study of Eight Countries. *Journal of Broadcasting & Electronic Media*, 68(5), 778–805. *Cited in Chapter 3: the eight-country survey evidence that prior deepfake exposure raises, rather than lowers, susceptibility.* https://doi.org/10.1080/08838151.2024.2410783.

Policy Clock: EU AI Act

[EU 2024] **European Union.** (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/20

24/1689/oj. *Cited in Appendix B as a policy clock for provenance demand; the regulation's dates only, no compliance claim.*

Law & Regulation

[SEC 2023] **U.S. Securities and Exchange Commission.** (2023, December 13). Standards for Covered Clearing Agencies for U.S. Treasury Securities and Application of the Broker-Dealer Customer Protection Rule With Respect to U.S. Treasury Securities. Release No. 34-99149; compliance dates subsequently extended (2025) to December 31, 2026 for eligible cash transactions and June 30, 2027 for eligible repurchase transactions. https://www.sec.gov/rules-regulations/2023/12/s7-23-22. *Cited in Section 3.1.3 for the phase-in of central clearing as a change to the balance sheet each unit of Treasury intermediation consumes—a move in the capacity constraint, not its removal. Dates as published at the time of writing; the reader should check the current compliance schedule.*

[SEC v. Howey 1946] **Securities and Exchange Commission v. W. J. Howey Co.**, 328 U.S. 293 (1946). https://supreme.justia.com/cases/federal/us/328/293/. *Source of the four-element investment-contract test — investment of money, in a common enterprise, with a reasonable expectation of profit, derived from the efforts of others — against which Section 27.2 reads the Value Capture Lemma's own design commitments. Cited for the test as stated by the Court; subsequent SEC staff statements and case law that apply it to digital assets are discussed in that section and are not reproduced here.*

[NIST 2026] **NIST (Moody, D., Perlner, R., Regenscheid, A., Robinson, A., & Cooper, D.)** (2026, October). *Transition to Post-Quantum Cryptography Standards* (NIST IR 8547, final). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.IR.8547. *Cited in Section 5.2 for the transition dates: quantum-vulnerable public-key algorithms, including ECDSA and EdDSA, deprecated after 2030 and disallowed under NIST standards and guidelines after 2035.*

Verification Asymmetry & PoUW (MatMul)

[Freivalds 1977] **Freivalds, R.** (1977). Probabilistic Machines Can Use Less Running Time. In *Information Processing 77: Proceedings of IFIP Congress 77*, pp. 839–842. Amsterdam: North-Holland. https://dblp.org/rec/conf/ifip/Freivalds77.html. *The $O(n^2)$ randomized check of a claimed matrix product $C = AB$ — pick a random vector v, compare $A(Bv)$ with Cv — that Section 22.2.1 takes as the canonical M2 (algebraic or probabilistic) verification and Section 22.3 builds the MATMUL acceptance test on. A theorem, not a conjecture: soundness $\leq 1/|\mathbb{F}|$ per round is proved. What remains conjectural (Komargodski et al., below) is the consensus-security property of a proof-of-useful-work built on it.*

[Chesney & Citron 2019] **Chesney, R., & Citron, D. K.** (2019). Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security. *California Law Review*, 107(6), 1753–1820. https://doi.org/10.15779/Z38RV0D15J. *Source of the term "liar's dividend": the benefit that accrues to a liar when the public knows fabrication is possible and so discounts authentic evidence. Cited in Chapter 3.*

[Fifth Circuit 2024] **United States Court of Appeals for the Fifth Circuit.** (2024). *Van Loon v. Department of the Treasury*, No. 23-50669 (5th Cir. Nov. 26, 2024). *Held that Tornado Cash's immutable smart contracts are not "property" within the meaning of IEEPA and therefore could not be designated by OFAC. OFAC delisted the contracts on 21 March 2025. Cited in Section 27.3 as the decision that defeated designation of the code.* https://www.ca5.uscourts.gov/opinions/pub/23/23-50669-CV0.pdf.

[SDNY 2025] **United States District Court, Southern District of New York.** (2025). *United States v. Storm*, No. 1:23-cr-00430 (S.D.N.Y.). Jury verdict, 6 August 2025. *Roman Storm convicted of conspiracy to operate an unlicensed money-transmitting business (18 U.S.C. §1960); jury did not reach a verdict on the money-laundering and sanctions-conspiracy counts. Cited in Section 27.3 as the decision that did not defeat prosecution of the coders. Read together with Van Loon: non-custodial design protected the protocol, not its developers.* https://www.courtlistener.com/docket/67720764/united-states-v-storm/.

[Rechtbank Oost-Brabant 2024] **Rechtbank Oost-Brabant.** (2024). Judgment in the prosecution of Alexey Pertsev (Tornado Cash), May 2024. *Conviction for money laundering in the Netherlands; cited in Section 27.3 as the first developer conviction in the Tornado Cash sequence. Under appeal at the time of writing.* https://uitspraken.rechtspraak.nl/details?id=ECLI:NL:RBOBR:2024:2069.

[OFAC 2022] **U.S. Department of the Treasury, Office of Foreign Assets Control (OFAC).** (2022, August 8). U.S. Treasury Sanctions Notorious Virtual Currency Mixer Tornado Cash. Press release. https://home.treasury.gov/news/press-releases/jy0916. *The designation action whose reach over immutable code the Fifth Circuit later rejected. Cited in Section 27.3.*

[OFAC 2025] **U.S. Department of the Treasury, Office of Foreign Assets Control (OFAC).** (2025, March 21). Tornado Cash Delisting. Press release. https://home.treasury.gov/news/press-releases/sb0057. *The delisting that followed* Van Loon*; read with the Storm verdict, which the delisting did not prevent.*

[FATF 2025] **Financial Action Task Force (FATF).** (2025). *International Standards on Combating Money Laundering and the Financing of Terrorism & Proliferation: The FATF Recommendations*, as amended; Recommendation 16 (wire transfers, the "Travel Rule") and its 2019 extension to virtual assets and virtual-asset service providers. Paris: FATF. https://www.fatf-gafi.org/en/topics/fatf-recommendations.html. *Cited in Chapter 27 for the Travel Rule obligations that bind intermediaries and not bearers.*

[EU 2023] **European Union.** (2023). *Regulation (EU) 2023/1114 of the European Parliament and of the Council of 31 May 2023 on markets in crypto-assets* (MiCA). *Official Journal of the European Union*, L 150, 9 June 2023. https://eur-lex.europa.eu/eli/reg/2023/1114/oj. *Cited in Chapter 27 for the issuer- and service-provider obligations against which the bearer posture is compared.*

[SEC 2025a] **U.S. Securities and Exchange Commission, Division of Corporation Finance.** (2025a). Statement on Certain Proof-of-Work Mining Activities. Staff statement, 20 March 2025. *Non-binding staff view that certain protocol mining activities on public, permissionless PoW networks do not involve the offer and sale of securities. Cited in Section 27.2; the statement addresses the activity of mining under a fixed protocol, not a token whose fee routing is governance-adjustable.* https://www.sec.gov/newsroom/speeches-statements/statement-certain-proof-work-mining-activities-032025.

[SEC 2025b] **U.S. Securities and Exchange Commission, Division of Corporation Finance.** (2025b). Statement on Certain Protocol Staking Activities. Staff statement, 29 May 2025. *Non-binding staff view that certain protocol staking activities do not involve the offer and sale of securities. Cited in Section 27.2 with the same limitation as the mining statement: it speaks to the activity, not to holder-steered economics.* https://www.sec.gov/newsroom/speeches-statements/statement-certain-protocol-staking-activities-052925.

[SEC 2026a] **U.S. Securities and Exchange Commission.** (2026a). Application of the Federal Securities Laws to Certain Types of Crypto Assets and Certain Transactions Involving Crypto Assets. Interpretive release, Release Nos. 33-11412 and 34-105020, File No. S7-2026-09; issued 17 March 2026, effective 23 March 2026. *Commission interpretation, joined by CFTC guidance. Cited in Section 27.2 for: the Howey test is not superseded (§I); governance tokens as digital commodities (§III.A); how a non-security crypto asset becomes subject to, and separates from, an investment contract, with liability for unperformed promises (§IV); protocol mining and staking (§V). Primary source for the section's frame; read 13 October 2026.* https://www.sec.gov/rules-regulations/2026/03/s7-2026-09.

[SEC Division of Corporation Finance 2026b] **U.S. Securities and Exchange Commission, Division of Corporation Finance.** (2026b). Frequently Asked Questions on Crypto Assets. Staff FAQ, issued 25 September 2026, updated 28 September 2026. *Staff views with no legal force, not approved by the Commission. Cited in Section 27.2 for Q2.5 on buyback programmes, protocol-funded burns included, which turns on the system being functional and having "no central party" (the latter added in the 28 September update).* https://www.sec.gov/about/divisions-offices/division-corporation-finance/faqs-crypto-assets.

[Komargodski & Weinstein 2025] **Komargodski, I., & Weinstein, O.** (2025). Proofs of Useful Work from Arbitrary Matrix Multiplication. *arXiv:2504.09971*. https://arxiv.org/abs/2504.09971. Submitted April 2025; last revised November 2025 (v4; versions 1–3 also list I. Schen).

Also in ESA 2026, LIPIcs 388, 2:1–2:17 (invited talk), https://doi.org/10.4230/LIPIcs.ESA.2026.2. [Note: the optimal-security claim is described as a *conjecture* in the abstract, not a proven guarantee. The paper supports "promising research direction," not "settled deployment."]

[Pass 2026] **Pass, R.** (2026). The Economics of Proof-of-Useful-Work. *arXiv:2606.06700*. https://arxiv.org/abs/2606.06700. *Cited in Section 13.3 for the equilibrium analysis of proof-of-useful-work in which the useful-work buyer's payment and the block subsidy compete for the same hardware, bearing directly on who pays for security when work has an external buyer. The thesis engages its regime analysis against its own fee-incidence result; no parameter is imported. The paper's first-page footnote states: "This work was done while consulting for Pearl Research Labs."*

[Basu 2026] **Basu, A.** (2026). The Usefulness Gap in Proof-of-Useful-Work: An Empirical Study of Pearl's cuPOW Protocol. *arXiv:2606.04819* (v2, preprint under review). https://arxiv.org/abs/2606.04819. *Cited in Table 22.5 for the deployment date only: a MatMul proof-of-useful-work mainnet launched in approximately March 2026. Its measurements of useful output, which the author qualifies as not conclusive, are not used.*

[Higham 2002] **Higham, N. J.** (2002). *Accuracy and Stability of Numerical Algorithms* (2nd ed.). Philadelphia: SIAM. https://doi.org/10.1137/1.9780898718027. *Source of the componentwise backward-error bounds for floating-point matrix multiplication — $|C - AB| \le \gamma_n |A||B|$ entrywise, with $\gamma_n = nu/(1-nu)$ for unit roundoff u — on which the relative, componentwise, operand-normalized form of the MATMUL acceptance bound in Section 22.3 rests. Cited for the bound's form; the acceptance tolerance is a publication choice pinned to a reference kernel, not a value taken from the book.*

Hardware Assurance & Supply Chain (Layer 0 Prior Art)

[Becker et al. 2013] **Becker, G. T., Regazzoni, F., Paar, C., & Burleson, W. P.** (2013). Stealthy Dopant-Level Hardware Trojans. In *Cryptographic Hardware and Embedded Systems — CHES 2013*, Lecture Notes in Computer Science, vol. 8086. Springer. https://doi.org/10.1007/978-3-642-40349-1_12. *Cited in Section 17.5.4 for the existence of hardware trojans implemented by changing only dopant polarity, which are invisible to optical inspection and to imaging-based reverse engineering. It is the reason the sampling program's detection bound carries a test-sensitivity term and is stated as a deterrent, not a bound, against this class; the bound is claimed only for counterfeit dies, binned-down parts, firmware substitution, and metering fraud.*

[Chuang et al. 2026] **Chuang, J., Seto, A., Berrios, N., van Schaik, S., Garman, C., & Genkin, D.** (2026). TEE.fail: Breaking Trusted Execution Environments via DDR5 Memory Bus Interposition. *IEEE Symposium on Security and Privacy*. https://tee.fail. *Cited in Appendix E for two findings stated on the project page: a DDR5 interposer setup costing under $1,000, and extraction of*

an Intel attestation key sufficient to produce TDX quotes that verify. No result about AMD attestation is used.

ZKML / Verifiable ML

[Peng et al. 2025] **Peng, Z., Wang, T., Zhao, C., et al.** (2025). A Survey of Zero-Knowledge Proof Based Verifiable Machine Learning. *arXiv:2502.18535*. https://arxiv.org/abs/2502.18535. *Cited in Section 20.7 for the state of the art in zero-knowledge inference — practical at small scale, a research problem at frontier scale. Survey; no result beyond that maturity characterization is imported.*

[Sun et al. 2024] **Sun, H., Li, J., & Zhang, H.** (2024). zkLLM: Zero Knowledge Proofs for Large Language Models. *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security*, 4405–4419. https://doi.org/10.1145/3658644.3670334. *Cited in Section 22.6 for one measurement (Table 1): a LLaMa-2-7B forward pass at sequence length 2,048 proved in 620 s on one NVIDIA A100 (40 GB), with a 2.36 s verifier. No production-readiness claim is drawn from it.*

FHE / FHE Rollups / OpenFHE / Fhenix

[Zyskind et al. 2024] **Zyskind, G., Erez, Y., Langer, T., Grossman, I., & Bondarevsky, L.** (2024). FHE-Rollups: Scaling Confidential Smart Contracts on Ethereum and Beyond. *Proceedings of BSCI '24*. ACM. https://doi.org/10.1145/3659463.3660031. *Landscape reference; not cited in the text.*

[OpenFHE 2025] **OpenFHE Project.** (2025). OpenFHE v1.4.0 – Development Release Announcement. GitHub. https://github.com/openfheorg/openfhe-development. *Landscape reference; not cited in the text.*

[Fhenix 2023] **Fhenix.** (2023). Introducing FHE-Rollups: Scaling Confidential Smart Contracts on Ethereum and Beyond. Fhenix blog. https://www.fhenix.io/blog/introducing-fhe-rollups-scaling-confidential-smart-contracts-on-ethereum-and-beyond. *Landscape reference; not cited in the text.*

Bitcoin Transport Privacy (BIP-324, Core 27.0)

[Bitcoin Core 2024] **Bitcoin Core Developers.** (2024, April 17). Bitcoin Core Version 27.0 Released. BitcoinCore.org. https://bitcoincore.org/en/releases/27.0/. *Cited in Chapter 18 for the release that enabled BIP-324 by default. A release note, cited as fact.*

[Mehta et al. 2023] **Mehta, D., Ruffing, T., Schnelli, J., & Wuille, P.** (2023). BIP-324: Version 2

P2P Encrypted Transport Protocol. Bitcoin Improvement Proposals. https://github.com/bitcoin/bips/blob/master/bip-0324.mediawiki. *Cited in Chapter 18, Appendix F, and Appendix A as the wire-format template for encrypted peer transport. A specification; nothing beyond its text is claimed.*

Lightning Receiver Privacy (BOLT12 Offers)

[NoBS Bitcoin 2024a] **No Bullshit Bitcoin.** (2024, July 4). Phoenix Wallet v2.3.1 & Phoenixd v0.2.0: BOLT12 Offers. NoBSBitcoin.com. https://www.nobsbitcoin.com/phoenix-wallet-v2-3-1-phoenixd-v0-2-0/. *Landscape reference; not cited in the text.*

[NoBS Bitcoin 2024b] **No Bullshit Bitcoin.** (2024, September 24). BOLT12: Offers Officially Merged into Lightning Specification. https://www.nobsbitcoin.com/bolt12-offers-officially-merged-into-lightning-spec/. *Cited in Chapter 18 and Appendix F for the date BOLT12 Offers merged into the Lightning specification. Secondary reporting of a specification event.*

Ecash / Federations (Fedimint / Fedi)

[Bitcoin Magazine 2025] **Bitcoin Magazine.** (2025, October 29). From Stealth to Scale: Fedi Unveils Multi-Sig Guardians for Federated Bitcoin Ecash Mints. https://bitcoinmagazine.com/business/from-stealth-to-scale-fedi-unveils-multi-sig-guardians-for-federated-bitcoin-e-cash-mints. *Landscape reference; not cited in the text.*

[Fedimint Project 2024–25] **Fedimint Project.** (2024–2025). What Does the Federation Do? Fedimint documentation. https://fedimint.org. *Cited in Chapter 18, Section 29.3, and Section 34.1.1 for what a federation does. Project documentation, read here as federated custody rather than in the project's own framing.*

Privacy Rails: Zcash, BTC↔XMR Swaps

[Hopwood et al. 2025] **Hopwood, D., Bowe, S., Hornby, T., & Wilcox, N.** *Zcash Protocol Specification* [NU6]. Electric Coin Company / Zcash Foundation. Rev. 2025.6.0-105. https://zips.z.cash/protocol/protocol.pdf. *Cited in Section 23.3 for what the shielded pool does and does not offer — in particular, no adaptor-compatible signing path — which is why the BTC↔shielded-ZEC corridor is gated. The protocol specification as the authority on the protocol.*

[Zcash 2026] **Zcash.** (2026, July 19). Orchard Action circuit soundness failure (under-constrained base in halo2_gadgets) enables counterfeiting within the Orchard pool. GitHub Security Advisory GHSA-ghc3-g8w4-whf9. https://github.com/zcash/zcash/security/advisories/G

HSA-ghc3-g8w4-whf9. *Cited in Section 30.1 for two statements in the advisory: the flaw "could have permitted counterfeiting (minting) of funds within the Orchard shielded pool," and "it is not possible to verify after the fact whether counterfeiting occurred during the exposure window." The advisory credits the ZIP 209 turnstile with bounding outflows from the pool.*

[Hopwood & Grigg 2025] **Hopwood, D.-E., & Grigg, J.** (2025). ZIP 2005: Ironwood Quantum Recoverability (status: Proposed). Zcash Improvement Proposals. https://zips.z.cash/zip-2005. *Cited in Section 5.2 for one statement: note encryption for the Zcash shielded pools is subject to "Harvest Now, Decrypt Later" attacks by an attacker holding the on-chain ciphertext and the receiver's address. A proposal, not a deployed rule.*

[Gugger 2020] **Gugger, J. (h4sh3d).** (2020). Bitcoin–Monero Cross-chain Atomic Swap. Protocol paper. IACR Cryptology ePrint Archive, Report 2020/1126. https://eprint.iacr.org/2020/1126. *The adaptor-signature construction for atomic swaps between a scriptable chain (Bitcoin) and a scriptless one (Monero) that Chapter 23's BTC↔XMR corridor takes as its operative reference. The protocol paper; the implementation is the following entry.*

[Rucknium 2025] **Rucknium.** (2025, September 26). Monero's 18-block reorg: What happened, why it matters, and what does it have to do with the annoying 10-block lock? https://rucknium.me/posts/monero-18-block-reorg/. *Cited in Section 23.4 for the event: an 18-block reorganization on 14 September 2025 that invalidated 115 transactions from the original chain. Analyst post; no claim about attribution is used.*

[COMIT 2021] **COMIT Network.** (2021). *xmr-btc-swap*: Bitcoin–Monero Atomic Swap Implementation. GitHub. https://github.com/comit-network/xmr-btc-swap. *The open-source implementation of the Gugger protocol, with its refund and punish transaction flows; the descendants listed below build on it.*

[UnstoppableSwap 2025] **UnstoppableSwap / eigenwallet.** (2025). UnstoppableSwap GUI / eigenwallet: Non-custodial XMR↔BTC Atomic Swaps. GitHub. https://github.com/eigenwallet/core. *Cited in Section 23.3 as a production descendant of the COMIT implementation. Existence and lineage only.* [Buterin et al. 2023] **Buterin, V., Illum, J., Nadler, M., Schär, F., & Soleimani, A.** (2023). Blockchain Privacy and Regulatory Compliance: Towards a Practical Equilibrium. *SSRN Working Paper*. https://doi.org/10.2139/ssrn.4563364. *Cited in Section 23.2 for association sets: a user can prove that their own inputs to a withdrawal belong to (or are excluded from) a published set without revealing which, which is why "none of my inputs is in exclusion set E" is provable while taint-at-distance is not.*

Zcash Roadmap / NU7 Discussions

[Zcash Forum 2024–25] **Zcash Community Forum.** (2024–2025). Accelerate NU7: Let's Deliver a Smaller Network Upgrade Sooner. Governance thread. https://forum.zcashcommunity.com/t/accelerate-nu7-let-s-deliver-a-smaller-network-upgrade-sooner/51916. *Landscape reference; not cited in the text.*

Ethereum Privacy L2 (Aztec)

[Steinrueck & Sammara 2025] **Steinrueck, D., & Sammara, A.** (2025, May 2). Testnet Retro – Launch Day: Aztec Public Testnet Went Live. Aztec Network blog. https://aztec.network/blog/testnet-retro---launch-day. *Landscape reference; not cited in the text.*

[CoinDesk 2025] **Coindesk.** (2025, May 1). "Everything Is Encrypted": Aztec's Privacy Rollup Hits Testnet Amid Growing Demand. https://www.coindesk.com/tech/2025/05/01/everything-is-encrypted-aztec-s-privacy-rollup-hits-testnet-amid-growing-demand. *Landscape reference; not cited in the text.*

IBC-Wide Shielding (Namada)

[Namada 2025] **Namada Team.** (2025, June 20). Namada Mainnet Launch Is Complete! Namada blog. https://namada.net/blog/namada-mainnet-launch-is-complete. *Landscape reference; not cited in the text.*

[Namada 2023] **Namada Team.** (2023, January 31). Understanding the MASP and CC Circuits. Namada blog. https://namada.net/blog/understanding-the-masp-and-cc-circuits. *Landscape reference; not cited in the text.*

Proof Markets / PoUW / Privacy-First L1s

[Aleo 2025] **Aleo.** (2025). Consensus and Proof of Succinct Work (PoSW). Aleo developer documentation. https://docs.aleo.org/learn/network/nodes/provers. *Landscape reference; not cited in the text.*

[Boundless 2025] **RISC Zero / Boundless.** (2025). Boundless: A Universal ZK Compute Layer. Binance Research report; the report page is not stably linkable, so the protocol documentation is given instead. https://docs.boundless.network. *Cited in Appendix E.6 for the market's auction structure. Documentation; no price is adopted from it.*

[Shankar 2026] **Shankar, S.** (2026, February 26). The cost of proving is now $0.02 per billion

cycles, an all time low since Mainnet launch in Q4. LinkedIn post (Boundless Network CEO). https://www.linkedin.com/posts/shiv-shankar_the-cost-of-proving-is-now-002-per-billion-activity-7432854966572244992-GIK-. *Cited in Appendix E.6 for the February 2026 Boundless auction reading (USD 0.02 per 10^9 cycles; roughly 400 trillion cycles per day of posted capacity). A vendor statement about a market price, cited as such; it is offered as evidence of oversupply in open proving auctions, not as a cost of proving.*

[Boundless 2026] **Boundless Foundation.** (2026, October 19). ZKC: The AI Access Coin. Blog post. https://boundless.network/blog/zkc-the-ai-access-coin/. *Cited in Appendix E.6 for one statement: "The Boundless ZK proving protocol will pause operations on November 5, 2026." The stated reasons are not used.*

[Succinct 2026] **Succinct Labs.** (2026). Succinct Prover Network documentation: Proof Contests; FAQ (pricing in prover gas units). https://docs.succinct.xyz/docs/protocol/spn/auction. *Cited in Appendix E.6 for the auction-set price per PGU plus dynamic base fee.*

[Succinct 2025] **Succinct Labs.** (2025, August 5). Succinct Prover Network Mainnet Launch (PROVE). https://blog.succinct.foundation/mainnet/. *Cited in Appendix E.6 for the mainnet launch date. Announcement only.*

[Nockchain 2025] **Nockchain Foundation.** (2025). Nockchain Is Now Open Source; Nockchain Issuance Schedule and Dumbnet Launch. Nockchain.org. https://www.nockchain.org/writings/nockchain-is-now-open-source. Issuance schedule: https://www.nockchain.org/writings/nockchain-issuance-schedule-and-updated-launch-target. *Cited in Chapter 1, Section 29.3, and Section 31.5 as a design pattern — zk-PoW with explicit issuance — and not as an endorsement of the project or its token; no quantity is adopted.*

[Psy Protocol 2025] **Psy Protocol (formerly QED).** (2025). Psy: A Scalable, Privacy-First L1 Secured by Proof-of-Useful-Work 2.0. https://psy.xyz/docs. *Landscape reference; not cited in the text.*

[Bittensor 2025] **Bittensor.** (2025). Yuma Consensus. Bittensor Docs. https://docs.bittensor.com. *Landscape reference; not cited in the text.*

[Ambient 2024] **Ambient.** (2024). Ambient Litepaper v1: Proof of Logits and Consensus. https://ambient.xyz/Ambient_Litepaper_V1.pdf.

Transport, Privacy, SDKs & Related Infrastructure

[NoBS Bitcoin 2023] **No Bullshit Bitcoin.** (2023, July 11). The New Phoenix: A 3rd Generation Lightning Wallet. https://www.nobsbitcoin.com/he-new-phoenix-a-3rd-generation-lightning-wallet/. *Landscape reference; not cited in the text.*

[Fedimint 2025] **Fedimint.** (2025). Fedimint Lightning Gateway Uses LDK Node to Simplify Deployment and Liquidity Management. Lightning Dev Kit blog. https://lightningdevkit.org/blog/fedimint-lightning-gateway-uses-ldk-node-to-simplify-deployment-and-liquidity-management. *Landscape reference; not cited in the text.*

Ethereum Privacy Roadmaps (PSE / Privacy Cluster)

[PSE 2025] **Privacy Stewards of Ethereum (PSE).** (2025, September 12). PSE Roadmap: 2025 and Beyond. PSE.dev blog. https://pse.dev/blog/pse-roadmap-2025. *Landscape reference; not cited in the text.*

[Ethereum Foundation 2025] **Ethereum Foundation.** (2025, October 8). The Ethereum Foundation's Commitment to Privacy. Ethereum.org blog. https://blog.ethereum.org/2025/10/08/privacy-commitment. *Landscape reference; not cited in the text.*

[CoinMarketCap 2025] **CoinMarketCap Academy.** (2025, October 9). Ethereum Foundation Launches Privacy Cluster for Enhanced Blockchain Security. https://coinmarketcap.com/academy/article/ethereum-foundation-launches-privacy-cluster-for-enhanced-blockchain-security. *Landscape reference; not cited in the text.*

Vitalik on Openness & Verifiable Hardware

[Buterin 2025] **Buterin, V.** (2025, September 24). The Importance of Full-Stack Openness and Verifiability. Personal blog (vitalik.eth.limo). https://vitalik.eth.limo/general/2025/09/24/openness_and_verifiability.html. *Landscape reference; not cited in the text.*

C2PA / Provenance (Supporting)

[C2PA spec] **Coalition for Content Provenance and Authenticity (C2PA).** *Content Credentials: C2PA Technical Specification*. A Joint Development Foundation project of the Linux Foundation; founding members Adobe, Arm, Intel, Microsoft, Truepic, and the BBC. https://c2pa.org.

Software Update Security & Transparency (Layer 2 Prior Art)

[TUF 2011–] **The Update Framework (TUF).** (2011–). *The Update Framework: A Framework for Securing Software Update Systems*. specification and reference implementations, theupdateframework.io. https://theupdateframework.io. *The supply-chain invariants I1–I5 in Chapter 19 (threshold signatures, freshness/anti-rollback, reproducibility, compromise response, multi-path avail-*

ability) restate, in protocol-specific form, guarantees TUF established for package repositories; Uptane (Kuppusamy et al., 2016–present) adapts the same framework to automotive update streams.

[Kuppusamy et al. 2016] **Kuppusamy, T. K., Brown, A., Awwad, S., McCoy, D., Bielawski, R., Mott, C., Lauzon, S., Weimerskirch, A., & Cappos, J.** (2016). Uptane: Securing Software Updates for Automobiles. In *Proceedings of the 14th Embedded Security in Cars Conference (ESCAR Europe)*. https://uptane.org. *The automotive adaptation of TUF: compromise-resilient update delivery to fleets of heterogeneous devices with an offline root, image and director repositories, and partial-verification clients. Cited as the closest prior deployment of the multi-role, multi-path update pipeline Chapter 19 specifies for a monetary client.*

[Laurie et al. 2013] **Laurie, B., Langley, A., & Kasper, E.** (2013). Certificate Transparency. RFC 6962, IETF. https://www.rfc-editor.org/rfc/rfc6962. *The append-only, publicly auditable release log of Section 19.3 is the CT construction applied to software artifacts; witness cosigning in the style of Sigsum (sigsum.org) is the lightweight variant.*

[Reproducible Builds 2014–] **Reproducible Builds Project.** (2014–). *Reproducible Builds: Setting the Record Straight* and associated specifications. reproducible-builds.org. https://reproducible-builds.org.

Censorship Measurement & Circumvention (Layer 1 Prior Art)

[OONI 2012–] **Open Observatory of Network Interference (OONI).** (2012–). *OONI Probe and Measurement Methodology*. ooni.org. https://ooni.org. *The VerifyReach measurement design (Chapter 18) — volunteer vantage points, canary requests, geographic/ASN stratification, public aggregation with p50/p95 reporting — follows the measurement methodology OONI has operated at scale for over a decade; the divergence triggers and WC-incentivized vantage pool are the extensions.*

[Dingledine et al. 2004] **Dingledine, R., Mathewson, N., & Syverson, P.** (2004). Tor: The Second-Generation Onion Router. *Proceedings of USENIX Security '04*. https://www.usenix.org/legacy/events/sec04/tech/dingledine.html. *Tor and I2P as first-class transports (Chapter 18) are prior deployments of the transport-agility and receiver-privacy principles; pluggable transports are the obfuscation layer referenced by "obfs-class shims."*

[Wu et al. 2023] **Wu, M., Sippe, J., Sivakumar, D., Burg, J., Anderson, P., Wang, X., Bock, K., Houmansadr, A., Levin, D., & Ensafi, R.** (2023). How the Great Firewall of China Detects and Blocks Fully Encrypted Traffic. In *Proceedings of the 32nd USENIX Security Symposium*. USENIX Association. https://www.usenix.org/conference/usenixsecurity23/presentation/wu-mingshi. PDF: https://www.usenix.org/system/files/usenixsecurity23-wu-mingshi.pdf. *Cited in Chapter 18 and Appendix F for the finding that the GFW has, since late 2021, blocked fully encrypted (uniform-random-bytes) transports heuristically — by entropy and printable-byte tests*

on the first packet — so that "looks like nothing" is itself a fingerprint. It is why the transport set separates encrypted transports from look-like-something transports and carries mimicry (look-like-TLS), Snowflake, and refraction (Conjure) fallbacks.

[Fanti et al. 2018] **Fanti, G., Venkatakrishnan, S. B., Bakshi, S., Denby, B., Bhargava, S., Miller, A., & Viswanath, P.** (2018). Dandelion++: Lightweight Cryptocurrency Networking with Formal Anonymity Guarantees. *Proceedings of the ACM on Measurement and Analysis of Computing Systems (SIGMETRICS)*, 2(2), Article 29. https://doi.org/10.1145/3224424. *Cited in Section 18.3 for two-phase transaction relay—an anonymity phase along a random privacy graph followed by ordinary diffusion—which bounds the first-spy estimator's advantage against mempool and gossip surveillance. Encrypted links do not address this threat; relay topology does.*

[Frolov et al. 2019] **Frolov, S., Wampler, J., Tan, S. C., Halderman, J. A., Borisov, N., & Wustrow, E.** (2019). Conjure: Summoning Proxies from Unused Address Space. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19)*. ACM. https://doi.org/10.1145/3319535.3363218. *Cited in Section 18.3 as the deployed form of refraction networking: a cooperating transit ISP redirects flows addressed to unused address space in its range, so that the censor sees a connection to a decoy and blocking it means blocking the ISP. Listed as a first-class look-like-something fallback beside Snowflake and domain fronting.*

[DoD 1989] **U.S. Department of Defense.** (1989). *MIL-STD-105E: Sampling Procedures and Tables for Inspection by Attributes*. Superseded for civilian use by **ANSI/ASQ Z1.4** (2003, reaffirmed 2018), *Sampling Procedures and Tables for Inspection by Attributes*. American Society for Quality. https://quicksearch.dla.mil/qsDocDetails.aspx?ident_number=35496. *Cited in Section 17.5.4 as the sampling precedent for the Layer 0 lot-inspection program: lot-by-lot attribute sampling with a published plan and a stated consumer's risk, in place of full inspection. Replaces the earlier and incorrect appeal to Lloyd's Register as a sampling precedent; Lloyd's classification was a per-vessel survey and is retained only as the precedent for a published, paid-for, third-party grade that underwriters price off.*

[Turan et al. 2018] **Turan, M. S., Barker, E., Kelsey, J., McKay, K. A., Baish, M. L., & Boyle, M.** (2018). *Recommendation for the Entropy Sources Used for Random Bit Generation* (NIST Special Publication 800-90B). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-90B. *Cited in Section 17.12 for the min-entropy estimators applied to a raw noise source, which is the correct test for an entropy source; statistical-randomness suites (SP 800-22, Dieharder, TestU01) test conditioned output and cannot certify a source. The Layer 0 RNG stress test applies 800-90B to the source and Dieharder/TestU01 to the conditioned output, not the reverse.*

Sourcing Note

This thesis cites primary documents, academic research, official specifications, and direct industry publications where possible, and every entry carries a URL or DOI: a DOI where one is issued; the publisher's, the court's, or the agency's canonical page otherwise; a library or Internet Archive record for books; and, for paywalled newsletters, the post's own URL with its publication date recovered from the publication's archive. Where an item has no page of its own (an aggregate data series; a report whose host is not stably linkable) the entry says so and links the nearest stable surface. Links were checked on 11 September 2026; a link that later rots is a defect in the record and is repaired in the next version, not silently removed. Where partisan reports or secondary journalism are cited, it is as a pointer to primary sources or as evidence of public discourse, not as an authoritative claim. Readers are encouraged to verify claims against the primary sources. IMF data has been updated to the April 2026 Fiscal Monitor where relevant.

Appendix A

Formal Model of Verification Asymmetry & VerifyPrice

A.1 Definition 1: VerifyPrice(W) and Verification Overhead

For each canonical workload W (for example, a constraint system $\text{PROOF}_{2^{20}}$, the matrix-multiplication SKU MATMUL_4096_FP32_FREIVALDS, or a verified-inference task INFER_LM_70B_256TOK two quantities matter:

1. How expensive it is to verify a purported result on a reference verifier, and
2. How that verification cost compares to producing the result from scratch.

VerifyPrice(W): verification profile Fix a reference verifier and adversarial mix of inputs. From a sample of verification runs for workload W, the following are estimated:

- $p_{50,t}(W), p_{95,t}(W)$: median and 95th-percentile wall-clock verify time from "artifact received" to "accept/reject";
- $p_{50,c}(W), p_{95,c}(W)$: corresponding median and 95th-percentile verify cost in a reference unit (e.g., USD of cloud compute);
- $\text{fail}(W)$: observed fraction of artifacts that fail verification (invalid, malformed, or timed out).

The VerifyPrice vector for W is then:

$$\text{VerifyPrice}(W) \equiv (p_{50,t}(W), p_{95,t}(W), p_{50,c}(W), p_{95,c}(W), \text{fail}(W))$$

Canonical form and its three views. This five-field tuple is the canonical definition of VerifyPrice, and every variant tuple elsewhere in the thesis is a view of it, not a competing definition. There are three, and they are related by a single derivation:

1. **Canonical (published) form:** $(p_{50,t}, p_{95,t}, p_{50,c}, p_{95,c}, \text{fail})$, as above. Every SLO, red line, and phase gate that names VerifyPrice reads on this form and no other (Section 22.7).
2. **Input view:** the per-run record $(t_{\text{run}}, e_{\text{run}}, m_{\text{run}}, b_{\text{run}}, \text{outcome}_{\text{run}})$ that the harness of Section 22.6 writes for every verification — wall-clock time, energy, peak memory, bytes transferred, and the pass/failure code — before any pricing or percentiling.

3. **Cost-vector view:** the same per-run record as Section 26.3 names it, CostVector$(W, V)_{\text{run}}$, and its front-matter form $c(W) = (t_{\text{cpu}}, m_{\text{peak}}, b_{\text{net}}, e_{\text{joules}}, c_{\text{est}})$, which adds the priced per-run cost c_{est} as a field.

The derivation from views (ii) and (iii) to the canonical form is *per run, then percentiled*: each run's cost is $c_{\text{run}} = (e_{\text{run}}/3.6 \times 10^6)\, p_{\text{energy}} + t_{\text{run}}\, p_{\text{opp}}$ with both reference prices pinned and republished on change (Section 22.6); $p_{50,c}$ and $p_{95,c}$ are then the percentiles of c_{run} across the window, never recomputed from percentile inputs. Only time and energy are priced. **Peak memory m and bytes transferred b are published, not priced**: they appear in receipts and in views (ii) and (iii), they bound the reference machine a verifier must be (Section 22.6), and they enter no cost percentile — a workload that fits the reference verifier's memory and bandwidth pays nothing extra for them in $p_{\cdot,c}$, and one that does not fit is not a canonical workload on that tier. A dashboard may show any view, but all three must reconcile to the same underlying receipts.

ProduceProfile(W): production profile The prover's cost to produce correct outputs for W is described the same way. From a sample of production runs (or cost models), estimate the median and 95th-percentile produce time and cost:

$$\text{ProduceProfile}(W) \equiv (p^{\text{prod}}_{50,t}(W), p^{\text{prod}}_{95,t}(W), p^{\text{prod}}_{50,c}(W), p^{\text{prod}}_{95,c}(W))$$

The denominator is an attack surface, and gets its own contract. The verification *numerator* of $r(W)$ is produced under the multi-operator, signed-batch, adversarial-corpus regime of Section 22.6; the production *denominator* has no equivalent protection as written, and it needs one, because a prover who inflates reported production cost drives $r(W)$ down and the network looks more verification-asymmetric than it is. The measurement contract for ProduceProfile(W) is therefore the mirror of the numerator's:

- **Reference production cost model per workload.** Production cost is estimated against a published reference model — pinned reference hardware, pinned kernel and stack versions, pinned energy price — exactly as the FP32 acceptance bound pins its reference kernel (Section 22.3). Self-reported production cost may be published but does not enter $r(W)$.
- **Independent estimation.** At least two independent estimators (distinct from the prover being measured) produce the reference-model estimate; divergence above 5% triggers the same investigation protocol as verification measurements.
- **Published derivation.** The mapping from reference hardware, energy price, and runtime to estimated production cost is public and reproducible, so any observer can re-derive the denominator from its inputs.
- **Asymmetry acknowledged.** Where independent estimation is impossible for a workload, $r(W)$ is published as *prover-favorable and unverified*, and cannot by itself support a passing reading on chain link 3 (Section 8.1).

The consequence for the SLO below is that a passing $r(W) \leq 0.3$ is a *joint* claim about both halves of the ratio: verification is cheap under the adversarial harness, *and* production cost is what the reference model says it is, not what the cheapest prover reports.

Verification overhead $r(W)$ Define separate time- and cost-based overhead ratios:

$$(v/p)_W^{\text{time}} = \frac{p_{50,t}(W)}{p_{50,t}^{\text{prod}}(W)}, \quad (v/p)_W^{\text{cost}} = \frac{p_{50,c}(W)}{p_{50,c}^{\text{prod}}(W)}$$

For convenience, these can be summarized as a single scalar:

$$r(W) \equiv \max\left((v/p)_W^{\text{time}}, (v/p)_W^{\text{cost}}\right)$$

A.1.1 Implementation Note: VerifyPrice Telemetry Schema (Sketch)

The exact encoding of VerifyPrice is implementation-dependent, but for operational dashboards and APIs it helps to standardize a minimal JSON-style schema per canonical workload. A skeletal example:

```
{
  "schema_version": "verifyprice/1",
  "workload_id": "PROOF_2^20",
  "modality": "M1",
  "window_start": "2026-06-01T00:00:00Z",
  "window_end":   "2026-06-07T00:00:00Z",

  "p50_verify_time_sec": 0.40,
  "p95_verify_time_sec": 0.90,

  "p50_cost_usd": 0.00030,
  "p95_cost_usd": 0.00055,

  "failure_rate": 0.0004,

  "r_time": 0.00035,
  "r_cost": 0.00050,
  "r_max": 0.00050,

  "sample_size": 125000,
  "hardware_profiles": ["open_tee_v1", "gpu_cluster_Y_v1"]
}
```

This schema is not normative; it is illustrative. The figures are chosen to be consistent with

the definitions: $\text{PROOF}_{2^{20}}$ is a succinct (M1) workload, so $r_{\max}$ sits inside the M1 band of 10^{-3} (Section 22.2.1) — a verify time of $0.4\,\text{s}$ against a production time on the order of 20 minutes gives $r_{\text{time}} \approx 3.5 \times 10^{-4}$, and the cost ratio is of the same order. A value of $r_{\max} \approx 0.2$ would place a SNARK workload two orders of magnitude outside its band and is the kind of reading the modality check exists to catch. The key properties are:

- Each canonical workload (workload_id) has its own series of VerifyPrice snapshots over sliding windows.
- VerifyPrice is expressed both in time and cost percentiles, with the overhead ratio $r(W)$ reported explicitly.
- Enough context is included to interpret shifts (sample size, hardware profile mix, time window), but no user-identifying data is required.

Clients and observatories can compute VerifyPrice from raw receipts and emit them in this or a similar format. The important thing is that:

1. The mapping from receipts $\rightarrow$ VerifyPrice is public and reproducible, and
2. The resulting metrics are exposed in a way that any third party can verify and alert on.

A system that claims "cheap, public verification" but cannot produce machine-readable VerifyPrice series for its canonical workloads is, by construction, asking you to trust what it cannot or will not measure.

W is called **verification-asymmetric** if $r(W) \ll 1$: checking honest work is much cheaper than re-doing it from scratch, both in time and cost.

A network (or facility) satisfies the VerifyPrice SLO for workload W if, for a published target (e.g., $p_{95,t}(W) < 5\text{s}$, $p_{95,c}(W) < 10^{-3}$ USD), it maintains:

- $\text{fail}(W)$ below a small, explicit threshold, and
- $r(W)$ within its modality band — 10^{-3} for succinct (M1), 10^{-2} for algebraic or probabilistic (M2) with the soundness parameter published, and no ratio target for replicated or attested verification (M3), which does not carry the monetary mechanism (Section 22.2.1). The legacy aggregate ceiling $r(W) \leq 0.3$ is retained as a portfolio-level engineering floor and as the fallback for unbanded workloads; no monetary claim reads on it,

over adversarial mixes of inputs and artifacts.

A.2 Definition 2: Facility Energy Receipt (FER)

A Facility Energy Receipt (FER) is a signed, machine-readable summary of a site's energy use and availability over a fixed time interval Δt. Each FER is identified by a unique hash `fer_id` and contains at least:

- **facility_id**: a stable identifier for the site or plant;
- **interval_start, interval_end**: timestamps delimiting the reporting window;
- **energy_in_kwh**: total electrical energy consumed at the site boundary over Δt;

- **energy_it_kwh**: energy delivered to IT loads (provers, routers, storage);
- **heat_reused_kwh**: energy captured and reused (for ERE metrics);
- **pue, ere, wue**: standard efficiency and water-use indicators over Δt;
- **outages**: a list of curtailment or outage events during Δt;
- **carbon_intensity_kg_per_kwh**: average emissions factor for supplied energy;
- **signature**: one or more digital signatures attesting to the report's correctness.

FERs are produced at regular cadence (e.g., every 5 minutes or hour). Other receipts (e.g., PIDL receipts for proofs) can reference a FER via `fer_id`, allowing allocators, auditors, and protocols to reconstruct the energy context in which a unit of verified work was performed.

A.3 Definition 3: PIDL Receipt (Proof Interface Definition Language)

A PIDL receipt is the minimal, canonical record of a verified interaction in the Create/Compute → Prove → Settle → Verify loop. It binds together the claim, the evidence, and the operational context into a single, serializable object.

At a minimum, a PIDL receipt contains:

- **receipt_id**: a unique identifier (e.g., a hash of the serialized payload);
- **claim_id**: an identifier or hash of the claim being attested;
- **workload_id**: a canonical registry identifier (e.g., PROOF_2^{20}, MATMUL_4096_FP32_FREIVALDS; Section 22.3.3);
- **proof_ref**: a hash or pointer to the proof or transcript artifact;
- **sla_tier**: the latency tier under which the work was promised (Bronze/Silver/Gold) and **assurance_tier**: the modality-indexed assurance tier (Replicated, Replicated+audit, Probabilistic, Succinct; Section 22.3.2) — a quote names one of each;
- **verify_time_sec, verify_cost_unit**: measured time and cost for verification;
- **result**: an outcome flag (e.g., accept, reject, timeout);
- **optional context**: hardware_profile, fer_id, settlement_txid, policy/jurisdiction tags, and principal_chain — an ordered list of commitments to the operator boundaries a delegated task crossed on its way to this receipt, each proven by predicate and none carrying an identity (Section 7.2, Section 26.8.9). The field is optional and additive: a receipt without it verifies exactly as before, and a verifier that does not understand it ignores it;
- **signatures**: digital signatures from the prover, verifier, and (if applicable) broker or router.

PIDL receipts are designed to be portable and composable: they can be embedded in blocks, stored off-chain with only their hashes committed on-chain, included in higher-level proofs, or presented to auditors and counterparties.

A.4 Definition 4: VerifyReach(N, R)

Reachability and path diversity for independent verifiers in network N and region R.

For a given network configuration N, client region R, and a target time budget T (e.g., "within 30 seconds"), VerifyReach summarizes how reliably an honest client can reach at least one honest verifier within T seconds over multiple transport classes.

The VerifyReach KPI for network N in region R is the vector:

$$\text{VerifyReach}(N, R) \equiv (p_{50,\text{conn}}(N, R), p_{95,\text{conn}}(N, R), \text{succ}_1(N, R), \text{succ}_2(N, R), \text{fail}_{\text{mix}}(N, R))$$

Where:

- $p_{50,\text{conn}}$, $p_{95,\text{conn}}$: median and 95th-percentile time-to-first-connection;
- succ_1: fraction of vantage points reaching an honest verifier via the *primary* transport within 30 s;
- succ_2: fraction reaching one via *any* transport, fallbacks included, within 60 s — so $\text{succ}_2 \geq \text{succ}_1$ by construction, and the gap between them is what the fallback ladder buys;
- fail_{mix}: distribution over connection failure causes.

SLOs, as Section 18.5 sets them (this definition carries no second set):

- $\text{succ}_1(N, R) \geq 0.95$ and $\text{succ}_2(N, R) \geq 0.99$ in uncensored regions; ≥ 0.70 and ≥ 0.85 in known-censored regions, measured over the ASNs sampled inside them;
- $p_{95,\text{conn}}(N, R)$ per transport class, because a single figure across classes hides the class that is failing: **direct encrypted** (BIP-324 [Mehta et al. 2023], QUIC/TLS) ≤ 3 s; **mimicry** (look-like-TLS, domain-front-free decoy routing) ≤ 8 s; **onion / I2P** ≤ 20 s; **rendezvous and refraction** (Snowflake-class WebRTC, Conjure-class [Frolov et al. 2019]) ≤ 45 s. Each is measured from first packet to an authenticated verifier session, and the class mix that produced the aggregate is published with it. The reference set is Appendix F.

The failure-cause distribution fail_{mix} is likewise reported per transport class, over a fixed taxonomy — DNS failure, TCP/QUIC reset, TLS/handshake failure, timeout with no response, active-probe detection, and application-layer block — so that "unreachable" is never published as a single number.

A.5 Definition 5: VerifySettle(C)

Time, reliability, and privacy quality of non-custodial settlement on a corridor C.

For a given settlement corridor C (e.g., BTC↔XMR atomic swaps), VerifySettle(C) summarizes whether non-custodial settlement actually clears under its advertised SLOs.

The VerifySettle KPI is:

$$\text{VerifySettle}(C) \equiv (p_{50,t}(C), p_{95,t}(C), s(C), r_{\text{safe}}(C), f_{\text{mix}}(C), a(C))$$

Where:

- $s(C)$ (success): fraction of settlements that complete as intended within the SLA window;
- $r_{\text{safe}}(C)$ (refund-safe): fraction of off-happy-path transactions in which every party had a unilateral no-loss exit within the corridor's bounded τ_C — its own leg refunded or the counterparty's leg claimed (Section 23.4); computed on the protocol-attributable denominator (user-environment failures classified and excluded) for the SLO, and on the overall denominator for Red Line 2 and the Phase I gate;
- $f_{\text{mix}}(C)$ (failure mix): distribution of failure causes;
- $a(C)$ (anonset): summary of anonymity-set size and churn.

Notation: $s(C)$ and $r_{\text{safe}}(C)$ are the mathematical symbols; where the body and the boards write `success(C)` and `refund_safe(C)` (Chapter 23, Red Line 2), those are the same two fields under their dashboard names, and nothing else is meant by either spelling.

A corridor satisfies its VerifySettle SLO if:

- $s(C) \geq$ target (e.g., 0.95), and
- $r_{\text{safe}}(C) \geq 0.999$ on the protocol-attributable denominator over adversarial mixes, with zero protocol-attributable losses (any such loss is Sev-1; Section 23.5), and ≥ 0.995 overall,

with no hidden custodian or discretionary freeze paths. A hard 1.0 over every failure class is a figure no corridor with humans on it can hold, and it is met only by corridors nobody uses.

A.6 Definition 6: Repression Wedge

The repression wedge is the realized real return on rate-capped instruments when nominal yields are held below inflation:

$$r_{\text{real},t} \equiv \frac{1 + i_{\text{cap},t}}{1 + \pi_t} - 1 \approx i_{\text{cap},t} - \pi_t$$

A negative $r_{\text{real},t}$ indicates that holders of capped instruments are being taxed in real terms via financial repression rather than explicit default.

A.7 Definition 7: Liquidation Effect

The liquidation effect measures the effective "tax" on bondholders from negative real rates. The construction follows Reinhart and Sbrancia [Reinhart & Sbrancia 2015], whose *annual liquidation effect* is the real saving to the sovereign in a year when the real rate on its domestic

debt is negative, expressed as a share of GDP:

$$\text{Liquidation}_t \approx \max(0, -r_{\text{real},t}) \times \frac{\text{Domestic Gov't Debt}_t}{\text{GDP}_t}$$

Reinhart and Sbrancia report that advanced economies saw negative real rates in roughly half of the years over 1945–1980, with the annual liquidation effect averaging on the order of 1–5% of GDP per year across the economies in their sample; those are their figures, quoted here as the historical scale of the mechanism, and no red line reads on them.

A.8 Definition 8: Physical VerifyPrice

Two series, two names. Two series are easily run together under one name. The constitutional SLO — the verification-cost series of Definition 1 measured in real resources on reference hardware, exogenous to token price, the series Red Line 1 reads — is the **Real-Resource VerifyPrice SLO** (Section 22.7). **Physical VerifyPrice** names only this definition: the cost of auditing the physical-infrastructure claims behind a unit of verified work, the series Red Line 10 reads. The two are not merged, because a network can verify proofs cheaply while its facilities are opaque, and the converse. Section 22.8 and the glossary carry the same distinction.

The claims being audited. A Facility Capacity Receipt (FCR, Section 17.7) is a signed, auditable receipt that extends the Facility Energy Receipt of Definition 2 with grid, cooling, redundancy, hardware-diversity, jurisdictional, and infrastructure-resilience claims; it is the object whose fields Sovereign Optionality and Delivered Verified Capacity aggregate. Physical VerifyPrice extends VerifyPrice(W) (Definition 1) from proof/compute claims to these FCR claims. For a facility profile F, it is the vector

$$\begin{aligned}\text{PhysicalVerifyPrice}(F) \equiv \big(&p^{\text{FCR}}_{50,t}(F),\ p^{\text{FCR}}_{95,t}(F),\ p^{\text{FCR}}_{50,c}(F),\ p^{\text{FCR}}_{95,c}(F),\\ &\text{ChallengeSuccess}(F),\ \text{OpenDiscrepancies}(F)\big)\end{aligned}$$

where $p^{\text{FCR}}_{\cdot,t}$ and $p^{\text{FCR}}_{\cdot,c}$ are time and cost percentiles to independently verify F's FCR claims (energy profile, outage/curtailment history, hardware-profile integrity, jurisdictional routing), $\text{ChallengeSuccess}(F)$ is the fraction of third-party audit challenges that resolve without dispute, and $\text{OpenDiscrepancies}(F)$ is the count of unresolved claim mismatches. A Work Credit is high-assurance only if VerifyPrice(W), PhysicalVerifyPrice(F), and the relevant DVC bound clear their respective SLOs; none of those conditions makes the credit monetary. See Sections 17.9 and 22.8.

A.9 Definition 9: VerifyFlow

VerifyFlow is the market-realization counterpart to VerifyPrice, VerifyReach, and VerifySettle. Where those three measure whether the stack supplies triad capacity under adversarial conditions, VerifyFlow measures the external financial representation and price-transmission state of the native monetary object. It is a pair of vectors, one per wrapper and one per asset:

$$\text{VerifyFlow}_i \equiv \big(L_i,\ A_i,\ \Delta A_i,\ \varepsilon_i,\ \kappa_i,\ C_i,\ H_i,\ D_i\big)$$

$$\text{VerifyFlow}(X) \equiv \big(Q^{RC},\ Q^{RD},\ \text{MPR},\ \text{CCR},\ \text{SER},\ \text{WNG},\ \text{RCR},\ \text{NUS}\big)$$

Component definitions appear in Section 26.8; the estimating equations and their measurement contract appear in Appendix H. VerifyFlow is explicitly *not* a protocol SLO: no VerifyFlow value falsifies the monetary thesis by itself. Its function is to determine whether observed price carries any information about native adoption (Corollary 13.1).

Appendix B

Practical KPIs & Telemetry Templates

This appendix collects the thesis's named instruments into one table — what each measures, in what unit, from what data, how often, on which public board it surfaces, which red line reads it, and at what threshold — and then lists the operational checklists by domain. The table is the reconciliation: a metric that appears in the body but not here, or here with a different definition, is a defect to report. Definitions are drawn from the body and the formal appendices (Appendix A, Appendix H); where a threshold is a red-line trigger it is the red line's own figure, and where the body publishes no number the cell says so rather than inventing one.

B.1 Instrument Table

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
VERIFYPRICE (W)	Canonical five-field tuple $(p_{50,t}, p_{95,t}, p_{50,c}, p_{95,c}, \text{fail})$ per canonical workload on the reference verifier under the adversarial input mix; cost derived per run from pinned energy and opportunity prices, then percentiled (Appendix A, Def. 1).	s; USD; fraction	Signed verifier receipts, multi-operator, aggregated by the observatory (Section 22.6)	Sliding window; daily, weekly published	Proof & Compute Board	RL1	Per-workload SLO (reference: SNARK $p_{95,t} \leq 5$ s, $p_{95,c} \leq \$0.01$, fail $\leq 0.1\%$); RL1 trips on p95 above SLO ≥ 3 consecutive months

Continued on next page…

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
$r(W)$, per modality	max of median time ratio and median cost ratio, verification over reference-model production (Appendix A); modality band M1/M2/M3 (Section 22.2.1).	ratio	VerifyPrice receipts; independently estimated reference production cost model	With the VerifyPrice window	Proof & Compute Board	None; link 3	M1 $\leq 10^{-3}$ (stretch 10^{-5}); M2 $\leq 10^{-2}$ with soundness published; M3 no target; legacy 0.3 is an engineering floor, no monetary claim reads on it
VerifyReach (N, R)	$(p_{50,\text{conn}}, p_{95,\text{conn}}, \text{succ}_1, \text{succ}_2, \text{fail}_{\text{mix}})$: time to first authenticated verifier session, one-path and multi-path reachability, failure-cause mix, per region and per transport class (Appendix A, Def. 4).	s; frac.; mix	Volunteer vantage probes paid from the assurance budget, OONI-style [OONI 2012–]; geo/ASN stratified (Chapter 18)	Live feed; daily publication	Settlement & Privacy Board; Layer 1 boards	None; L1 falsifier	$\text{succ}_1 \geq 0.95$ / $\text{succ}_2 \geq 0.99$ uncensored; ≥ 0.70 / ≥ 0.85 known-censored (Section 18.5); p_{95} per transport class (Appendix A); L1 falsified if $> 20\%$ of sampled ASNs *in uncensored regions* unreachable > 7 days

Continued on next page...

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
VERIFYSETTLE (C)	$(p_{50,t}, p_{95,t}, s, r_{\text{safe}}, f_{\text{mix}}, a)$: settlement time, success, refund-safety, failure mix, anonymity-set summary per corridor (Appendix A, Def. 5).	s; frac.; mix	Corridor receipts, refund transaction IDs, per-asset anonymity metrics (Chapter 23)	Live feed; weekly per corridor	Settlement & Privacy Board	RL2	$s \geq 0.95$; $r_{\text{safe}} \geq 0.999$ on the protocol-attributable denominator with zero protocol-attributable losses (Sev-1 on any); RL2 trips on any protocol-attributable refund failure on an admissible route, or overall $r_{\text{safe}} < 0.995$ with failures unclassified, repeated and with no automatic delist
VerifyFlow	Per-wrapper vector $(L, A, \Delta A, \varepsilon, \kappa, C, H, D)$ and asset vector $(Q^{RC}, Q^{RD}, \text{MPR}^{\text{gross/net}}, \text{CCR}, \text{SER}, \text{WNG}, \text{RCR}, \text{NUS})$ (Section 26.8, Appendix A.9).	mixed	Fund shares outstanding, NAV, AUM, creations; venue volumes; custodian-tagged balances; derivative open interest and disclosures	Daily or intraday; rolling windows	Market Realization & Wrapper Board	None (warnings)	No VerifyFlow value falsifies the thesis by itself; Market Realization Warnings trip on elasticity break, κ sign change, MPR above depth bounds, creations dominating native spot

Continued on next page…

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
Wrapper Dominance Ratio (WDR)	Pair: stock $E^{\text{wrapper}}/(E^{\text{wrapper}} + E^{\text{native}})$; flow $V^{\text{wrapper}}_{t-h:t}/V^{\text{native}}_{t-h:t}$ with h published; never divided into one scalar (Section 13.6).	fraction; ratio	ETF and wrapper AUM, custodied and exchange balances, wrapped supply; native fee, settlement, and collateral series	Quarterly	Wrapper Dominance Board	RL9	Both components rise ≥ 2 consecutive quarters while native fee share, private settlement volume, and collateral lockups stagnate or decline; split published if they diverge
Delivered Verified Capacity (DVC)	Scenario-specific maximum surviving flow from energy and fuel through firm power, switchgear, cooling, hardware, network, workload, proof, verification, settlement, to usable service; with active binding constraints, substitution latency, common-cause groups (Section 17.9).	service units per period	FCR and FER fields; dependency graph; published disruption scenarios	Per scenario; recomputed quarterly and on scenario change	Layer 0 Capacity Board; Common-Cause Dependency Board	RL15 (stress DVC); RL13 upstream	Issuance capped by stress-adjusted DVC, not nameplate; zero-DVC scenarios excluded only by explicit force-majeure declaration; RL15 clauses conjunctive
FeeCoverage	Retained fees as a share of the security budget, burns excluded (Section 25.6.6); the take rate τ (gross fees over turnover) is a separate series.	fraction	Chain data: fee routing, issuance, burn ledger	Monthly, trailing window	Value Capture Board; Economic Coverage Board	RL5	RL5: $< 10\%$ with workload mix $> 80\%$ speculative for ≥ 12 months; 30% sustained 12–24 months is the II→III phase gate, 50% the III→IV gate (Section 29.3)

Continued on next page…

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
Homestead Ratio	Share of verified compute, proof generation, and AI-service capacity supplied by open-admission, non-hyperscaler, geographically diverse, independently verifiable operators (Section 5.6).	fraction	VerifyPrice-tracked capacity by operator class, jurisdiction, and hardware profile	Quarterly	Layer 0 Capacity Board; Value Capture Board (triple divergence)	RL11 (via complement)	Enclosure Risk Flag / RL11: top-3 hyperscalers, closed TEEs, or a single jurisdiction $> 50\%$ of tracked capacity, persisting; 40% chip-family/TEE cap is the early-warning tier
Mechanical Pressure Ratio (MPR)	Gross $\sum_j \lvert Q_j^{\text{mkt}}\rvert / V_X$ (pressure); net $\lvert Q^{\text{mkt}}\rvert / V_X$ (direction); never averaged or combined (Section 26.8.2).	ratio (pair)	Q^{mkt} from wrapper flows via dealer-hedge decomposition; V_X from reference-venue volumes	Daily; same window numerator and denominator	Market Realization & Wrapper Board	None (Warning 3)	Pre-declared market-depth bounds, published with the ratio; the reporting partition (venue $\times$ wrapper class) published alongside
Capital Survival Ratio (CSR)	Whether successive investor cohorts in a wrapper preserved capital, as distinct from whether the product preserved AUM (Section 26.8.4).	frac. of cohorts	Shares outstanding and NAV history by launch vintage, chained across corporate actions	Quarterly	Market Realization & Wrapper Board	None (RL9 support)	No numeric trigger; published per cohort so "AUM grew" cannot be read as "holders were made whole"

Continued on next page…

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
Sovereign Optionality ($\mathcal{O}_s$, $\mathcal{O}_s^{\text{net}}$)	Capacity-weighted composite of substitutability attributes from FCR fields; network figure is the capacity-weighted mean less an HHI concentration penalty over jurisdictions and balancing authorities (Section 17.8).	index	Facility Capacity Receipts (Section 17.7)	Quarterly	Layer 0 Capacity Board	RL13	$\mathcal{O}_s^{\text{net}}$ below published threshold ≥ 2 consecutive quarters, or $> 50\%$ of tracked capacity in jurisdictions operating active curtailment or rationing against verification workloads
Physical VerifyPrice (FCR audit cost)	Time and cost to independently verify a facility's FCR claims, plus challenge-success rate and open discrepancies (Appendix A, Def. 8; Section 22.8).	s; USD; frac.; count	Third-party FCR audit challenges and their resolutions	Quarterly	Layer 0 Capacity Board; Economic Coverage Board	RL10	Capacity-weighted p_{95} time or cost exceeds $2\times$ the frozen trailing-four-quarter baseline for ≥ 2 consecutive quarters, or FCR data unavailable or unverifiable for $> 20\%$ of active capacity for ≥ 2 consecutive quarters; watch indicator before the Phase II gate

Continued on next page…

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
Off-ledger exposure disclosure	Share of pledged-position notional whose operator publishes an attested disclosure of off-ledger leverage, counterparty concentration, and rehypothecation — the scope-drift gap of Section 13.3.10: a position legible on its proof and opaque on its operator's balance sheet.	fraction	Operator attestations bound to PIDL receipts; lender panel of RL17	Quarterly	Value Capture Board	None (feeds RL10 and RL17 readings)	No numeric trigger in this version; published so the gap between what is verified and what is exposed is a number rather than a caveat
M1+M2 share of verified units	Fraction of the network's verified output whose correctness anyone can check under a succinct (M1) or algebraic/probabilistic (M2) modality; the share carrying the monetary mechanism at all (Section 22.2.1).	fraction	Receipts by registry SKU and modality	With the VerifyPrice window	Value Capture Board	Invariant (a) Sev-1; RL3 second reading	Falling below the published floor: Sev-1; a migration of the mix from M1 toward M3 is the quiet form of RL3's monoculture

Continued on next page…

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
Convenience-yield vector (RL16)	Base-asset lending rates by venue, term, and borrower class; forward and perpetual basis net of the reference rate; wrapper basis against native holdings; each read against the regime-pressure index R, with the corridor band fixed by the pre-committed classifier and independent access readings (Section 26.8.6).	bp; corridor increment θ_{16} in bp	Lending venues, derivatives venues, wrapper NAV; borrow demand decomposed by borrower class first	Daily series; quarterly estimate	Market Realization & Wrapper Board	RL16	Upper bound on the corridor increment below the pre-registered minimum over ≥ 8 quarters containing ≥ 2 transitions, kills only at $\geq 80\%$ pre-registered power; *inconclusive* while the interval admits the minimum; *insufficient market* below Phase III scale

Continued on next page…

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
Haircut dispersion (RL17)	Cross-lender interquartile range of quoted haircuts at a common timestamp on (i) the native asset, (ii) pledged protocol positions, (iii) an opaque comparable set (wrapped/custodial representations, tokenized commodity claims), residualized on trailing realized volatility and quoted depth (Section 13.3.10).	percentage points	Pre-committed panel of ≥ 5 lenders under RL14 custody and reproducibility terms	Quarterly; continuous during declared stress	Value Capture Board	RL17	Native dispersion across independently maintained underwriting approaches not lower than the comparables' over ≥ 4 quarters, or widening relative to them (or re-converging more slowly) in a declared stress episode while receipts and VerifyPrice stay within SLO; *insufficient panel* below $N = 5$; *unidentifiable* below five independently maintained approaches

Continued on next page…

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
Free-choice denomination share and quote/settle matrix (RL18)	Fraction of contracts the protocol did not require to be priced in the unit whose price terms are written in it, by count, notional, tenor, and counterparty tier (I/II/III), subsidy-adjusted, with the four-cell quoting-unit $\times$ settling-unit matrix and redenomination events at renewal; never aggregated with the mandated complement (Section 13.7.3).	fraction; notional	Contract registry classified by a party other than the fee recipient; accounting carve-out applied first	Quarterly	Value Capture Board	RL18	A (weakens): Tier II/III share flat or declining over 4 quarters while throughput and fee turnover grow, fiat-quoted natively-settled cell holding share. B (kills): insiders' unit-denominated share of eligible short-tenor obligations (price fixed ≤ 90 days) below the $\geq 25\%$ floor at Phase III entry, or below it for 4 consecutive quarters after; sales on receipt are evidence, not a trigger; long-tenor share published, no fatal threshold. Gate: $\geq 10\%$ of turnover free-choice, ≥ 50 counterparties

Continued on next page…

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
Acceptance breadth	Number of independent venues, corridors, and counterparties that quote, settle, and cite this instance rather than any other — the observable form of Δ_{coord} (Section 13.3.8); the lowest-fee protocol-grade fork is the mandatory fork row of the RL14 reference panel.	count	Venue registry; RL14 reference panel	Quarterly	Value Capture Board	None; RL14 input	No trigger of its own; the fork row binds RL14 Condition A first because it carries no cost handicap
Forced disclosure incidence; non-custodial share (RL12)	Share of surveyed employment, banking, housing, travel, and benefits onboarding events requiring a viewing key or full-history disclosure; share of settled native volume completed without a custodial or KYC chokepoint (Section 27.1.2).	fraction	Published survey panel; chain data classified under public criteria	Quarterly	Agency Preservation Board	RL12	Incidence $\geq 10\%$ in any major jurisdiction, or non-custodial share $< 50\%$ while institutional usage grows, either for ≥ 12 months
Governance-weight concentration; emergency-path use (RL8)	Largest address family's share of time-locked governance weight (Section 25.5); count of emergency-path invocations (Section 25.5.3) over a trailing four-quarter window.	fraction; count	Chain data (lock contracts, signer log)	Per vote; quarterly	Neutrality & Admission Board	RL8	$> 20\%$ for ≥ 2 consecutive quarters; > 2 invocations in 4 quarters; structural clauses (constraints absent, scope eroded) are binary

Continued on next page…

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
Native Monetary Buyer Map	Just-in-time fee acquisition; operator sell-through and inventory; gross burns, gross and net issuance; self-custodied reserve accumulation under published criteria; wrapper demand; leverage; holding-period distributions; loss-bearing capacity; countercyclical accumulation (Section 26.8.8), each with its weighting rule disclosed.	mixed	Chain cohorts classified conservatively; custodian-tagged balances; derivative disclosures	Quarterly	Native Monetary Buyer Map	None; link 8 and Premise 7	No numeric trigger; Premise 7 passes only when the retained, self-custodied, unlevered or transparently financed, loss-bearing, countercyclical cohort is material and persistent
Principal Map (overlay)	Principal shares of native fee turnover, prover collateral, corridor liquidity, and the self-custodied reserve cohort by class (human with standing, firm, agent policy, unknown); Human-Revocable Share (HRS); Native Policy Concentration Ratio (NPCR) with exit correlation; agent-originated share by operator class (Section 26.8.9).	share	Predicate-proven control proofs; policy template hashes; operator attestations; residue published as unknown	Quarterly	Overlay on Native Monetary Buyer Map, Value Capture, Agency Preservation, Layer 0 Capacity	RL12 watch clause; candidate RL19	Placeholders only: HRS $<$ [50%] for ≥ 4 quarters with rising agent share (watch); NPCR $>$ [50%] or top operator class $>$ [70%] for ≥ 2 quarters (candidate, weakens)

Continued on next page…

Metric	Definition	Unit	Data source	Cadence	Public board	Red line	Threshold
Cross-Perimeter Share (CPS)	Fraction of receipt-bearing delegation hops whose `principal_chain` crosses an operator boundary (Section 7.2).	share	PIDL receipts carrying the optional `principal_chain` field; settlement-graph inference as a labeled weaker fallback	Quarterly; *not computable* until field coverage is material	Layer 0 Capacity Board	None; second reading of RL1 (Corollary 30.1)	No numeric trigger; falling CPS with VERIFYPRICE inside SLO is the signature of enclosure by convenience
Capturable wedge (ω) and take rate (τ)	$\omega = s - \gamma$: normalized bypass spread less verification cost share; τ = gross native fees over settled protocol turnover, compared against ω (Chapter 30, RL14).	ratio	Reference panel of bypass prices (with mandatory fork row); chain fee and turnover data; workload unit-price deflator	Quarterly	Value Capture Board	RL14	Condition A on τ against ω; Condition B on real fee turnover failing to grow while physical throughput grows, with persistence and the named deflator (Chapter 30)

B.2 Operational Checklists by Domain

The table above holds the named instruments. Operators also need the ordinary telemetry beneath them, grouped here by domain. None of these carries a red line by itself; each feeds one of the instruments above or one of the twelve boards (Section 26.9).

Prover markets: Queue depth, clearing price per proof type, cancellation rates, p50/p95 VerifyPrice.

Compute consensus pilots: MatMul throughput (GFLOPs/s), verifier cost ratio $r(W)$ with its modality band, invalid-work rejection rate under adversarial load, adversarial acceptance-suite pass rate (Section 22.3).

Privacy settlement: Atomic swap success %, time-to-finality, stuck-flow causes by f_{mix} category, wallet UX friction metrics.

Decentralization & capture risk:
- Nakamoto coeff. (by stake/hash/prover share); top-N concentration; Gini.
- Entry latency: time-to-first-proof/mined-block for a new node.
- Geographic/ASN diversity; % over Tor/VPN; relay/MEV-builder diversity.

Liveness & safety:
- Reorg/orphan rate; effective finality time (p50/p95); incident MTTR/MTBF.
- SLA attainment (% proofs/settlements meeting latency/uptime targets).
- DoS/churn resilience: success % under adversarial load mixes.

Verify/produce economics:
- Verifier cost ratio (verify/produce) by workload; energy per proof/inference.
- Hardware mix & utilization (GPU/CPU/ASIC) vs. throughput; cost/FLOP ($/GFLOP-verified).
- Failure taxonomy: prover faults, invalid proofs, circuit timeouts.

Privacy rails health:
- Anonymity-set size & churn (shielded pool/epoch); linkability regressions found.
- Swap slippage & fees (p50/p95); cross-venue route diversity; failure-mix root causes, classified under the fixed f_{mix} taxonomy of Appendix A, Definition 5, rather than free text.
- UX friction: successful first-run rate, steps/clicks to complete, abandonment rate.

Bridges & cross-chain settlement:
- Atomicity violations (0-conf leaks, partial completes); retry rate; liquidity depth per route.
- Watchtower/guardian coverage; light-client verification share vs. multisig.

Provenance & policy demand:
- C2PA [C2PA spec] attach rate, strip rate by platform; "verified views" share.
- Proof volume indexed to policy milestones (EU AI Act [EU 2024] dates); % budgeted vs. ad-hoc spend.

Security & abuse:
- Sybil detection hit rate; staking/miner collusion alerts; MEV/censorship events.
- Key/attestation hygiene: validator/prover key rotations, slashing/penalties.

Token & issuance telemetry:

- Realized issuance vs. schedule; FeeCoverage (retained fees, burns excluded) of security budget; % rewards tied to useful work.
- Velocity vs. locked/collateralized supply; SoV signals (holder age bands, exchange outflows).

Customer success:

- Enterprise proof spend: $/month, SLA breaches, churn; **Net Proof Retention (NPR)** — for the cohort of customers that were paying for proofs in a base period, their proof spend in the current period divided by their base-period spend, computed on receipt-tagged budgets and excluding new customers; NPR > 1 means the existing base grew its verification demand net of churn and down-sell, the proof-market analogue of net revenue retention.
- Time-to-integration (SDK→prod); support tickets per 1k proofs/settlements.

Hardware honesty:

- Share of receipts tied to open or sampled hardware profiles (by workload and volume).
- Lot-sampling coverage: percentage of lots and devices subject to destructive audits (SEM/optical imaging, side-channel tests), and maximum time since the last audit per profile.
- Incidents and deprecations: count and severity of profile-level compromises, time-to-deprecation, and migration progress away from affected profiles.

Appendix C

The SDK (Proofs-as-a-Library)

At the top of the stack, the application layer, you work with a small vocabulary:

- "Prove this computation."
- "Prove this provenance."
- "Settle this payment under these privacy and finality constraints."
- "Anchor this result to hardware with this security profile."

The SDK turns those sentences into types. A claim is a first-class object: ComputationClaim, ProvenanceClaim, SettlementClaim, optionally decorated with requirements. The compiler and runtime then map those claims to whatever combination of proving backends, settlement rails, and verifiable machines currently clear the SLA at the best VerifyPrice.

Core types:

```
class Policy:
    max_verify_time_sec: float
    max_verify_cost_usd: float
    privacy_level: str              # e.g. "encrypted_io", "public"
    finality_target: str            # e.g. "1m", "30m", "2h"
    allowed_hardware_profiles: list # e.g. ["open_tee_v1", "pure_zk_only"]

class ComputationClaim:
    fn: Callable
    inputs: Any
    policy: Policy

class ProvenanceClaim:
    media_bytes: bytes
    device_profile: str
    policy: Policy

class SettlementClaim:
    payments: list  # [Payment(recipient, amount, asset), ...]
    policy: Policy

class Receipt:                      # = PIDL receipt, Def. 3 (app:verifyprice)
```

```
    receipt_id: str                  # hash of the serialized payload
    claim_id: str                    # hash of the claim being attested
    workload_id: str                 # canonical tag, e.g. "PROOF_2^20"
    proof_ref: str                   # hash/pointer to proof or transcript
    sla_tier: str                    # latency: "bronze" | "silver" | "gold"
    assurance_tier: str              # "replicated" | "replicated_audit" |
    "probabilistic" | "succinct"
    verify_time_sec: float
    verify_cost_unit: float          # in the published reference unit
    result: str                      # "accept" | "reject" | "timeout"
    hardware_profile: str | None
    fer_id: str | None               # Facility Energy Receipt, if bound
    settlement_txid: str | None
    policy_tags: dict                # jurisdiction / policy context
    signatures: list                 # prover, verifier, [broker/router]
    schema_version: str              # "pidl/1"
```

The `Receipt` type is the PIDL receipt of Appendix A, Definition 3, field for field; the SDK adds no fields the receipt does not carry, so a receipt exported from one application verifies in another and aggregates into the same VerifyPrice series (Section 22.6). `proof_bytes` travels alongside the receipt, addressed by `proof_ref`, rather than inside it.

SDK surface:

```
sdk = ProofSDK(networks=[...])  # PoUW chains, zk rollups, open-TEE clusters

claim   = sdk.make_claim(...)
receipt = sdk.prove_and_settle(claim, pay_with=user_wallet)
result  = sdk.verify(receipt)   # -> {accept: bool, metadata: ...}
```

Everything else—choice of proving backend, choice of settlement rail, choice of hardware profile—is a routing decision made under the hood, constrained only by the policy you declared.

Error model. Every call either returns a `Receipt` whose `result` is `accept`, or raises. Failures are typed, not silent: `PolicyUnsatisfiable` (no backend clears the declared `Policy` at the current VerifyPrice — the SDK never quietly relaxes a policy to find a route); `ProofRejected` (a receipt with `result = reject`, returned inside the exception so the rejection itself is evidence); `VerifyTimeout` (`result = timeout`, counted in fail(W)); and, for settlement claims, `SettlementAborted`, which carries the refund receipt — a settlement that fails must end in `refund_safe`, and an abort path that cannot produce a refund receipt is a Red Line 2 incident, not an SDK error. Rejected and timed-out receipts are emitted to telemetry exactly as accepted

ones are; a client that drops them corrupts the failure-rate field of every series it feeds.

Versioning. Receipts, policies, and hardware-profile identifiers are versioned independently and each carries its own tag: `schema_version` on receipts (`pidl/1`), a `policy_version` on policies, and the trailing `_vN` on profile strings (Appendix E). The SDK follows semantic versioning; a change to receipt fields or to the meaning of a policy field is a major version. A verifier must reject a receipt whose schema version it does not recognize rather than verify the subset it understands, and observatories publish VerifyPrice series per schema version until a migration completes, so a schema change is never mistaken for a cost change.

Example 1: Proofed medical inference on verifiable machines

```
policy = Policy(
    max_verify_time_sec = 1.0,
    max_verify_cost_usd = 0.001,
    privacy_level = "encrypted_io",
    finality_target = "30s",
    allowed_hardware_profiles = ["open_tee_v1", "pure_zk_only"]
)

@proofed(policy=policy)
def diagnose(image: EncryptedImage) -> Diagnosis:
    return model.predict(image)

# In request handler:
claim   = sdk.make_computation_claim(fn=diagnose, inputs=enc_image, policy=policy)
receipt = sdk.prove_and_settle(claim, pay_with=clinic_wallet)
result  = sdk.verify(receipt)
```

Example 2: Camera provenance anchored in open hardware

```
# On capture (running on "open_camera_v1" device):
raw_bytes = camera.capture()

policy = Policy(
    max_verify_time_sec = 0.5,
    max_verify_cost_usd = 0.0005,
    privacy_level = "hide_location_and_identity",
    finality_target = "5m",
    allowed_hardware_profiles = ["open_camera_v1"]
)

prov_claim = sdk.make_provenance_claim(
```

```
    media_bytes=raw_bytes,
    device_profile="open_camera_v1",
    policy=policy
)

receipt = sdk.prove_and_settle(prov_claim, pay_with=newsroom_wallet)
bundle = MediaBundle(media=raw_bytes, provenance_receipt=receipt.export())
```

Example 3: Private payroll over neutral rails

```
policy = Policy(
    max_verify_time_sec = 3.0,
    max_verify_cost_usd = 0.002,
    privacy_level = "shielded_settlement_with_auditable_receipts",
    finality_target = "2h",
    allowed_hardware_profiles = ["pure_zk_only"]
)

# Amounts are in the base asset (placeholder symbol NATIVE). A stablecoin
# leg here would be a Value-Capture Condition 1 bypass (part-ii,
# lem:value-capture) and appears in none of the reference applications.
payroll_batch = [
    Payment(recipient=alice_addr, amount=1000, asset="NATIVE"),
    Payment(recipient=bob_addr,   amount=1200, asset="NATIVE"),
]

settle_claim = sdk.make_settlement_claim(payments=payroll_batch, policy=policy)
receipt = sdk.prove_and_settle(settle_claim, pay_with=treasury_wallet)
```

The asset choice in that batch is not incidental. The Value Capture Lemma's first condition (Lemma 13.1) is that users cannot pay in fiat, stablecoins, or other tokens at equivalent service quality, and Section 13.5.1 names dollar stablecoins as the incumbent bypass channel; a reference application that settled payroll in one would be demonstrating the failure mode the thesis is instrumented to detect. `NATIVE` is a placeholder for whatever the base asset's ticker becomes; the reference applications carry no other settlement asset.

Appendix D

Energy & Plant Architecture

Plant architecture for verifiable work Objective: deliver verifiable, dispatchable, and composable power to proving and inference clusters so that proof units and verified FLOPs remain cheap to check and neutral to route.

Electrical topology (MV → LV):

- Medium-voltage interconnect with dual utility feeds where available; 2N or N+1 step-down to rack power.
- UPS/BESS for ride-through and proof-preserving shutdowns.
- Power-quality SLOs (voltage/frequency/THD) bound worst-case verifier latency inflation.

Thermal topology:

- Air to liquid (direct-to-chip/immersion) as default for high-density provers.
- Heat-reuse loops to district heating, greenhouses, or absorption chillers to improve ERE.
- Publish PUE and ERE: if heat is monetized, ERE < 1.0 is both possible and economically material.

Network & provenance:

- Redundant backhaul with clock discipline (PTP/Stratum) so receipts carry tight timing jitter bounds.
- Facility attestation path: meters and controllers sign energy/thermal telemetry.

Control plane:

- Workload-aware orchestration couples grid signals → job scheduler.
- Interruptible Bronze proofs pause first; Gold proofs ride through with BESS and on-site firming.

Energy procurement & grid posture:

- Portfolio, not a point bet: long-dated PPAs, nodal index exposure, behind-the-meter firming, and BESS.

- Co-location with stranded or curtailed energy (hydro spill, wind/solar curtailment, flare-gas).
- Demand response as a feature: treat prover/inference farms as dispatchable loads.
- Carbon & provenance: bind certificates into FER format so "kgCO_2e per proof" is auditable.

SLA tiers that map power quality to receipts:

- **Bronze (interruptible):** Shed priority 1. Eligible for demand-response; refunds/penalties via SLA escrow if p95 VerifyPrice breached. Typical use: PoL audits, non-urgent prover workloads.
- **Silver (curtailable with notice):** Shed priority 2 with N+1 firming; curtailment windows posted ahead of time. Typical use: batch MatMul-PoUW, rollup proving backlogs.
- **Gold (firm):** 2N electrical path or N+1 + BESS ride-through; zero curtailment except force majeure. Typical use: settlement-critical proofs, receipt ledger finalization.

What a lab would need to specify. The architecture above is a pattern, not a plant. Nothing in this appendix carries a budget, and none should be read into it; the numbers belong to a site study, not a thesis. Before a reference facility can be costed, the following have to be fixed, and the investment addendum (Chapter 32) is where the first of them are:

- **MW scale.** The IT load the reference plant is designed for, stated as a range with the prover and verifier shares separated, since verification is threatened by hardware obtainability and proving by bulk power (Section 17.10.1) and the two do not scale together.
- **PUE and ERE targets.** Design-point and worst-season figures for each, with the heat-reuse path that makes ERE < 1.0 achievable named or the target set above 1.0; both are FER fields (Appendix A.2) and the targets are what the receipts will be read against.
- **BESS sizing.** Energy (MWh) and power (MW) capacity, ride-through duration at Gold load, cycling assumptions, and whether the battery is sized for proof-preserving shutdown only or for firming as well.
- **PPA tenor and structure.** Contract length, fixed versus indexed pricing, curtailment terms, and the share of load covered by contract versus nodal exposure; the portfolio posture above becomes a set of instruments only once tenors are chosen.
- **Mapping of SLA tiers to FER/FCR fields.** For each of Bronze, Silver, and Gold: which `outages` entries count as curtailment against the tier, the interconnection class and backup-duration fields of the FCR (Section 17.7) a facility must report to offer the tier, and the SLA-escrow rule that reads them, so that "Gold" is a field-level claim a receipt can falsify rather than a marketing label.
- **Grid and jurisdiction.** Balancing authority, interconnection queue position, and the curtailment or rationing posture of the host jurisdiction toward verification workloads, since the last is what Red Line 13 reads.

Appendix E

Hardware Profiles

This thesis treats "hardware profiles" as first-class objects: they describe machines in the same language used for proofs and receipts.

E.1 Concept

A hardware profile is a public description of a class of devices that are interchangeable for security and performance purposes. It is not a single physical unit; it is the specification for what counts as a valid unit of that type.

Informally: a profile says, "devices that meet this description behave like this, within these error bars, under this sampling regime."

Hardware profiles serve to:

- express trust and threat-model assumptions at the protocol and application layer;
- bind proofs and receipts to the types of machines that produced them; and
- make the honesty (and failure rates) of machines measurable over time.

E.2 Naming

Profiles are identified by short, stable strings. Three broad kinds:

1. **Open-hardware profiles** — devices whose RTL/microarchitecture and sampling regime are publicly documented.
   - `open_tee_v1` — first open TEE design with published RTL, side-channel budget, and lot-sampling plan. No such device yet exists; the profile string names the specification a lab would build to (below).
   - `open_signer_v1` — simple key-storage/signing chip with open RTL and sampling.
   - `open_camera_v1` — camera/capture pipeline with documented sensor path, secure element, and sampling method.
2. **Logical or pure-cryptography profiles** — paths that do not rely on special hardware.
   - `pure_zk_only` — verification/proving path relying only on general-purpose CPUs/GPUs and cryptographic assumptions.

- `software_only_v1` — profile used for testing or where hardware attestations are unavailable.

3. **Vendor or mixed profiles** — existing devices not fully open but characterized and sampled.
   - `vendor_tee_X_v1` — proprietary TEE profile with documented side-channel and attestation scheme, and the vendor's stated scope for physical attacks; memory-bus interposers costing under $1,000 have extracted attestation keys and forged attestations [Chuang et al. 2026].
   - `gpu_cluster_Y_v1` — particular GPU farm configuration with known performance and sampling behavior.

E.3 Hardware Profile Specification

Each profile has a specification that is part of the public "bill of materials" for the stack. At a minimum:

- **profile_id**: string (e.g., "open_tee_v1")
- **device_class**: what the device is used for (e.g., "TEE", "signer", "accelerator", "camera")
- **trust_mode**: "open", "mixed", "opaque", or "logical"
- **design_commit**: identifiers for design artifacts (hashes of RTL, GDS, firmware images)
- **side_channel_budget**: qualitative/quantitative bounds
- **sampling_plan**: how lots and devices are checked
- **attestation_scheme**: how devices prove identity and state
- **security_target**: high-level security goal

E.4 Hardware Profiles in Receipts and Policies

Profiles appear in three main places:

1. **Policies** — Developers express requirements using profiles:

```
allowed_hardware_profiles = ["open_tee_v1", "pure_zk_only"]
```

2. **Receipts (PIDL)** — Every PIDL receipt includes a hardware_profile field.
3. **Telemetry** — The observability layer aggregates metrics by profile.

E.5 Example Profiles

open_tee_v1:

- device_class: "TEE"
- trust_mode: "open"

- design_commit: hashes of RTL, layout, firmware
- sampling_plan: random lot sampling, SEM/optical inspection
- intended use: high-value proving, key storage, secure computation

open_camera_v1:
- device_class: "camera"
- trust_mode: "open"
- sampling_plan: device sampling and destructive testing per lot
- intended use: origin-anchored capture for media provenance

pure_zk_only:
- device_class: "logical"
- trust_mode: "logical" (no hardware assumptions)
- intended use: contexts where hardware trust is unacceptable

E.6 What a Lab Would Need to Specify

The profiles above are specifications for describing machines. They are not machines, and this appendix attaches no prices to any of them; a bill of materials with costs belongs to a procurement study against the ladder of Section 17.4, not to a thesis. What can be stated here is the *shape* of what a lab would have to specify before the first receipt could carry an open profile string honestly.

L0-A reference prover and verifier: component list, no prices. The reference verifier is the machine every VerifyPrice figure is measured on (Laptop tier, Section 22.6), so its bill of materials is a public artifact even though every part in it is commodity. The reference prover is the machine the production denominator of $r(W)$ is priced against (Appendix A).

- **Reference verifier:** commodity x86-64 or ARM host meeting the Laptop-tier specification (4 cores, 16 GB, SSD, 100 Mbps); a pinned operating-system image and kernel version; the reference verifier software build, reproducible from source (Chapter 19); an external power meter at the wall for the energy field; a hardware RNG source for the adversarial-mix sampler; and a signing key on a commodity secure element for result attestation. Nothing in it is open silicon: L0-A is the *characterized commodity* rung, and the profile says so.
- **Reference prover:** commodity GPU or CPU cluster node with the pinned kernel and driver stack that sets the FP32 acceptance bound (Section 22.3); metered rack power distribution for the per-run energy figure; a pinned firmware inventory (hashes published); and the same reproducible software build.

- **For both:** the design_commit is the set of hashes of every image the machine runs; the sampling_plan is electrical test and firmware-hash comparison on a random sample of units per lot, since counterfeit dies, binned-down parts, and firmware substitution are the high-sensitivity defect classes (Section 17.5.4).

Proving-cost roadmap: the metric, today's reading, and the bet. The lab's headline cost series is stated in a unit a buyer can price: US dollars per 10^9 zkVM cycles for general programs, and per 10^9 constraints for circuit-native workloads, measured on the reference prover above as *cost*—hardware amortization plus metered energy, in the per-run price model of Section 22.6—and published beside the market *price* on open proving markets, which is a different number and must not be substituted for it. The distinction is load-bearing at the outset because the market reading today sits below any plausible cost floor. The open proof markets that launched in 2025 (Boundless [Boundless 2025]; the Succinct Prover Network [Succinct 2025]) clear reverse auctions in which posted prover supply exceeds order flow: Boundless reported [Shankar 2026] roughly USD 0.02 per 10^9 cycles in February 2026 against some 400 trillion cycles per day of posted capacity, with auction clears near zero at times, and Succinct [Succinct 2026] prices in prover gas units at an auction-set rate plus a dynamic base fee. A price set by oversupplied provers competing for subsidised order flow is evidence about the auction, not about the cost of proving, and the lab's series is the cost. Boundless announced in October 2026 that its proving protocol would pause on 5 November [Boundless 2026], which is consistent with that reading. This paragraph deliberately fixes no figure for that cost; the first quarterly reading does. The author's prior, stated so it can be wrong, is that all-in cost on commodity GPUs sits one to two orders of magnitude above the February 2026 auction clear, which is what oversupply predicts.

Two series ride beside the headline. The first is the **proving overhead** $p(W)/p_0(W)$: the production cost of a result *with* its proof over the cost $p_0(W)$ of executing the same workload natively with no proof, per registered SKU. It is the number the buyer actually pays for verifiability and the one the capturable wedge (Section 22.10) is measured against; a workload's M1 SKU becomes economic when that overhead falls inside the buyer's differential Δ for it, and not before. The second is the same overhead for the M3 alternative—a committee of k re-executing the work, whose overhead is roughly k and falls only as hardware does.

The **modality bet** is then a statement about slopes. The monetary mechanism rides on M1 and M2 (Section 22.2.1); the bet is that M1 proving overhead per unit of useful work falls faster than the M3 committee's, so that the set of workloads for which succinct verification is economic widens over time rather than staying fixed at the small-circuit end. Its checkpoints are dated to the lab's deliverables (Table 32.1): first published reading at $T_0 + 6$; first annual slope at $T_0 + 12$; and a pre-registered kill condition at eight quarters—if the measured decline in M1 overhead over that window is no steeper than the decline in the M3 committee's, the bet

is lost for that workload class, the class stays M3 in the registry (Table 22.4), and the affected `_ZKML`-style destination SKUs remain targets rather than becoming constitutional. A projection with no such condition would be a forecast, and this appendix publishes instruments.

The first L0-C open components. Open PDKs exist today only at mature nodes, so the first open components are the ones whose job is trust rather than throughput (Section 17.4). A lab would specify, in this order:

- **Root of trust (RoT):** an open RISC-V measured-boot and key-storage core, with published RTL, a documented attestation scheme, and a side-channel budget; this is what `open_signer_v1` names.
- **Random number generator (RNG):** an open entropy source and conditioner with a published health-test specification [Turan et al. 2018], since every sampling plan and every adversarial mix in the thesis draws on it.
- **Metering:** an open power and capacity meter whose firmware signs the FER fields (Appendix A.2), so that "energy_it_kwh" is attested by a component the operator does not control alone.

Side-channel budget metric. A profile's `side_channel_budget` is stated in a form an independent lab can reproduce: Test Vector Leakage Assessment (TVLA), as the pass/fail leakage test at a stated trace count and statistical threshold, with the ISO/IEC 17825 test procedure as the reference method. The published budget names the trace count at which the device is claimed to pass, and a device is re-tested at that count on every sampled lot.

Sampling plan reference. Lot acceptance follows a published acceptance-sampling standard — ANSI/ASQ Z1.4 [DoD 1989] (the successor to MIL-STD-105) — with the acceptable quality level, inspection level, and switching rules stated in the profile; the rule-of-three converse and the sensitivity term of Section 17.5.4 then convert the plan into a published upper bound on undetected compromise for the high-sensitivity defect classes, and into a deterrent statement, not a bound, for dopant-level trojans [Becker et al. 2013].

Appendix F

Communications Resilience Mechanisms

P2P transport set:

- v2 Encrypted P2P (BIP-324 [Mehta et al. 2023]) as default for all full nodes and relays; expose QUIC/TLS fallbacks.
- Onion/I2P modes first-class in clients; auto-failover if clearnet handshakes exhibit DPI resets.
- Handshake camouflage for relays in high-interdiction ASNs — with the caveat that fully encrypted, uniform-random-bytes transports (Noise, obfs4-class) are themselves a fingerprint the Great Firewall has blocked heuristically since 2021 [Wu et al. 2023], so camouflage must include *look-like-something* transports, not only *look-like-nothing* ones.

What a lab would need to specify: the transport list. The transport set is stated as classes so that VerifyReach can publish p_{95} per class (Appendix A, Definition 4), and each class names a reference implementation a lab would integrate rather than write:

- **Direct encrypted:** BIP-324 [Mehta et al. 2023] v2 transport (reference: Bitcoin Core [Bitcoin Core 2024]); QUIC/TLS 1.3 with Encrypted Client Hello pinned.
- **Fully encrypted obfuscation:** obfs4-class pluggable transports (reference: Tor pluggable transports). Retained for jurisdictions that do not run entropy-based blocking; published as its own class because it fails first where they do.
- **Mimicry (look-like-TLS):** transports whose wire image is a valid TLS session to a plausible host — uTLS-style fingerprint mimicry and WebTunnel-class HTTPS tunnels (reference: Tor WebTunnel) — so the flow is indistinguishable from ordinary web traffic to a passive observer.
- **Rendezvous:** Snowflake-class WebRTC through ephemeral volunteer proxies, with a domain-fronted or AMP-cache broker (reference: Tor Snowflake), for the case where every fixed endpoint is blocked.
- **Refraction:** Conjure-class refraction networking, in which the client connects to an unused address in a partner ISP's space and a station in the path redirects the flow (reference: Refraction Networking / Conjure [Frolov et al. 2019]); the highest-latency and highest-partnership-cost class, and the fallback of last resort.
- **Onion routing:** Tor and I2P as first-class transports, not fallbacks, for receiver privacy as

well as reachability [Dingledine et al. 2004].

The client's failover order is a published policy, not a heuristic, and every fallback taken is recorded in fail_{mix} so that a region quietly living on refraction is visible on the board.

Receiver-private routing for payments:

- BOLT12 Offers [NoBS Bitcoin 2024b] (receiver-private, reusable invoices) in wallets and merchant stacks.
- Path-blinding where supported.
- Federated ecash [Fedimint Project 2024–25] and shielded pools as optional "last-mile" sinks/-sources.

Settlement survivability:

- Adaptor-signature atomic swaps [Gugger 2020] for BTC↔XMR corridors with typed failure & refund flows; the BTC↔shielded-ZEC corridor is gated on shielded-pool signing work that does not yet exist (Section 23.3).
- Wallet UX exposes "abort & refund" as a first-class path.

Topology & anti-eclipse hardening:

- Peer-set diversity targets (geo/ASN spread), randomized peer rotation, inbound slot quotas.
- Gossip-path multiplicity to reduce single-jurisdiction capture.
- Light-client modes that verify receipts over any reachable path (browser, mobile, enclave).

Appendix G

Glossary of Terms & Notation

Triad Privacy, Proofs, and Compute — three cryptographic service capacities that may support a monetary candidate. Whether they should share one asset is tested by the Triad Coherence Test, not assumed.

Triad Coherence Test The test of whether the three triad services should share one asset at all. Privacy settlement, portable proof, and verified compute are different goods with different users, bottlenecks, collateral needs, political risks, and duration warehouses; a common asset is justified only if shared security, settlement, liquidity, and transaction-cost benefits exceed cross-subsidy, governance, and wrong-way-risk costs. Three architectures are under test — (A) one native monetary asset, (B) a neutral reserve plus service-specific credits, (C) shared settlement plus modular domain collateral — and none is pre-selected; Architecture A is a candidate, not the thesis. See Section 8.3.

Conditional chain The nine-link argument by which the thesis establishes its claim rather than asserting it: utility demand $\rightarrow$ standardized work $\rightarrow$ cheap-to-check receipts $\rightarrow$ stress-deliverable service $\rightarrow$ markets $\rightarrow$ fee flows $\rightarrow$ a scarce asset (only under non-bypassable accrual) $\rightarrow$ a persistent, self-custodied, loss-bearing, regime-responsive holder constituency $\rightarrow$ a possible store-of-value premium. Every arrow is a conditional; each link names where it is defended and what would sever it, and Table 8.1 rates them. Red Lines 1–15 read on links 2–8; Red Lines 16–18 read on link 9. See Section 8.1.

VerifyPrice(W) A public KPI vector measuring the cost and time to verify workload W. Canonical form $(p_{50,t}, p_{95,t}, p_{50,c}, p_{95,c}, \text{fail})$; the input and cost-vector tuples elsewhere are views of it (Appendix A).

VerifyReach(N,R) Metrics for network reachability under censorship conditions.

VerifySettle(C) Metrics for settlement success and safety on corridor C.

Value Capture Lemma The bridge from service use to asset accrual in the reference design: demand for triad capacity reaches the native asset through five channels the design relies on together—required fee medium, supply reduction via burns, collateral lockup, issuance discipline, and non-bypassability. Each is necessary to this design's route, not to value capture in general. It establishes a cash-flow claim and a collateral lockup, not monetary premium. See Lemma 13.1.

$r(W)$ Verification overhead ratio: the maximum of the median time ratio and the median

cost ratio, $r(W) = \max\big(p_{50,t}(W)/p^{\text{prod}}_{50,t}(W),\ p_{50,c}(W)/p^{\text{prod}}_{50,c}(W)\big)$, verification over production. Read against a modality band (M1: 10^{-3}; M2: 10^{-2}; M3: no ratio target), with the legacy 0.3 ceiling retained as an engineering floor on which no monetary claim reads. See Appendix A, Section 22.2.1.

VerifyPower Token-quoted VerifyPrice, a market signal and not a constitutional target: VerifyPower(token; W $q_{\text{fee}}(W^*) \cdot P_X / p_{50,c}(W^*)$ for a single pinned reference workload W^*, where q_{fee} is the fee schedule in units of the native asset and P_X its USD reference price, so the token price appears explicitly on the right-hand side. A vector has no reciprocal; the scalar is constructed by naming the workload and the field. The token-side image of the AffordabilityRatio, published together with it. See Section 22.7.

AffordabilityRatio VerifyCost(W)/MedianFee(W), the operational SLO that verification stays affordable as a fraction of typical transaction cost; target $\leq 5\%$ for all workload classes, with a settlement variant (verification cost $\leq 1\%$ of median USD-referenced transaction value). Its perverse direction — improving when fees rise — is held in check by three guards: component disclosure, USD-referenced units for the settlement variant, and a monotonicity rule crediting only improvements achieved on a flat or falling fee. Operational, not constitutional. See Section 22.7.

Base Asset The protocol's bearer fee and settlement unit and the only instrument tested as a monetary candidate. Fees, burns, collateral, or governance use establish neither moneyness nor a durable holder anchor by themselves.

Work Receipt (WR) A PIDL artifact proving that a specific unit of work was completed under attested conditions: claim hash, proof hash, workload ID, SLA tier, timestamps, hardware profile, prover signature. Copyable, verifiable by anyone, *not scarce*. It proves the past and confers no future rights; Work Receipts are evidence, not money; verified delivery against a Work Credit issues one and extinguishes the credit. See Section 12.1.

Work Credit A typed, prepaid claim on reserved future service, owed by a named provider or pool: issued against an admissible capacity commitment with sale and reservation atomic, extinguished by verified delivery, which issues a Work Receipt. Outstanding claims are bounded by stress-adjusted deliverable capacity. It may be workload-, location-, hardware-, SLA-, and time-specific; it is not presumed to be a store of value. See Section 12.1.1.

Bypass Risk The risk that users consume triad capacity (proofs, privacy, compute) through channels that do not require the native asset—e.g., fiat-denominated cloud services, stablecoin payment to operators. If bypass channels dominate, the SoV thesis fails.

Differential value (Δ) The marginal buyer's willingness to pay for the protocol's differential properties—cryptographic checkability, neutrality, credible non-discretion, non-custodial exit—over the best bypass channel. Bounds the sustainable fee: $f \leq \Delta - (c_p - c_b)$, where c_p and c_b are the protocol's and the bypass channel's all-in costs. Also written $\Delta(R)$ when its dependence on regime pressure is at issue. A scalar wedge in price units, not the difference

operator. See Section 13.3.3.

Fee incidence The question of who bears a protocol fee. With free entry and roughly constant returns, long-run supply is horizontal at marginal cost, so the fee is passed to buyers and operator economic profit is unchanged: a fee is a wedge on *turnover*, not a claim on residual profit. With installed capacity fixed, part of the fee is instead absorbed as quasi-rent to deployed hardware, which affects entry. See Section 13.3.1.

Capturable wedge (ω) $\omega = s - \gamma$, where s is the normalized bypass spread $(P^{\text{prot}} - P^{\text{byp}})/P^{\text{byp}}$ and γ is the verification cost share (proving plus redundancy) as a fraction of P^{byp}. The measurable proxy for Δ: what remains of the price the protocol commands after paying for the verification that justifies it. Normalized because the absolute spread shrinks mechanically with compute deflation. Read by Red Line 14, Condition A.

Realized take rate (τ) *Gross* native fees in USD divided by settled protocol turnover in USD — gross, not net of the burned share, because τ measures what buyers pay; burns are already inside gross fees and adding them would count destruction as payment (Section 25.6.6, Red Line 14). Computable by any third party from chain data. Compared against ω rather than against a guessed threshold, which is what makes Red Line 14's first condition self-normalizing.

Regime-Contingent Convenience Yield The channel through which a triad asset could earn monetary premium: the *holder-side service flow* arising from the state-contingency of $\Delta(R)$—the bearer's ability to transact, prove, hold, and exit when the substitutes for proofs, courts, and custodians have stopped working—plus credible non-discretion and bearer holdability. It accrues by virtue of holding rather than as a distribution, so a valuation of distributions omits it; a valuation of possession need not, and a benefit of ownership is not yet monetary use. Against Bitcoin converted on the day it reduces to a balance that stays spendable when acquisition fails (see Access Experiment). Red Line 16 reads its corridor increment. Named by analogy to the convenience-yield *mechanism* of Kaldor (1939) and Working (1949), not by any claim that a triad asset is a storage commodity. Explicitly *not* the countercyclicality of the fee stream, which is a negative-beta discount-rate effect a DCF captures exactly; and explicitly *not* produced by burns, collateral, or fee share, which yield a cash-flow claim and a collateral lockup. Unsized. See Section 13.3.11.

Collateral unit-elasticity The property that a collateral requirement specified as a share of capacity value locks a *value* rather than a quantity: required units are $N = \chi K/P$ for collateral ratio χ, capacity value K, and price P. Consequences: a level effect rather than a growth effect, a bound of χK set by an accounting identity, wrong-way risk in large drawdowns, and an effect on the liquidity state Λ_t rather than the monetary impulse θ_t. See Section 13.3.12.

Utility-Token Trap The pattern where a network provides useful services but the native token fails to capture economic value because operators extract all surplus, users pay in alternative currencies, or governance inflates supply. The thesis's value-capture conditions are designed

to prevent this outcome.

Information-insensitivity The property that no counterparty gains from acquiring private information about the particular unit or position being tendered, so it is accepted at face value without due diligence. The working definition of money-likeness in the collateral and short-term-funding literature. Scoped throughout this thesis to *unit quality* rather than to fundamentals: an information-insensitive asset can be violently volatile. See Section 13.3.10 and Sources [Dang et al. 2020].

No questions asked (NQA) Acceptance of an asset at par without investigation of the specific unit [Holmström 2015]. The observable consequence of information-insensitivity, and what Red Line 17 tests through haircut dispersion. See Section 13.3.10.

Symmetric knowledge The route to no-questions-asked acceptance in which producing information about the collateral's unit quality costs approximately nothing *and the result is public*, so no counterparty can acquire an informational edge by producing it — whether or not anyone does. The mechanism is deterrence on the claimant, not audit by the counterparty: a misrepresentation any party could falsify at zero cost is not worth attempting, so verification is constantly available and almost never used, and the object to measure is its cost, not its rate. Contrasted with *symmetric ignorance*, the historical route, in which producing that information is prohibitively expensive so nobody does — which holds only until somebody does, that being the run. See Section 13.3.10.

Haircut dispersion The cross-lender spread in required over-collateralization against the same asset at a common timestamp. Preferred over haircut *level* because it nets out price volatility; read relatively — native asset against an opaque comparable quoted by the same lenders, across independently maintained underwriting approaches — because lenders also differ in funding, balance sheet, and terms, and a shared model can produce agreement without knowledge. Primary series for Red Line 17. See Section 13.3.10.

Free-choice denomination share The fraction of contracts, among those the protocol did *not* require to be denominated in the base asset, whose price terms are nonetheless written in it. Reported by count, notional, tenor and counterparty tier, and never aggregated with the mandated complement, which is a routing rule rather than evidence. Primary series for Red Line 18. See Section 13.7.3.

Verification modality (M1/M2/M3) The three ways of establishing that work was done correctly, differing by orders of magnitude and in kind: succinct proof (M1; target $r(W) \leq 10^{-3}$, stretch 10^{-5}), algebraic or probabilistic check with published soundness error (M2; target $r(W) \leq 10^{-2}$, Freivalds-style matrix verification [Freivalds 1977] the canonical case), and replicated or attested execution (M3; no ratio target, because the ratio is not the operative quantity). Only M1 and M2 deliver the symmetric knowledge the monetary mechanism requires; M3 supports service delivery, Work Credits and SLA tiers but not the pledgeability claim, because a position resting on a committee requires a view about that committee. See Section 22.2.1.

Stage A index (R_t) The regime-pressure index of the three-stage demand model: R_t indexes financial repression, surveillance, synthetic media, and AI concentration, and generates demand for triad capacity, $D_t^{\text{triad}} = g(R_t)$. Compute demand is deliberately *not* a component of R_t — it is a demand driver in its own right, entering alongside regime pressure — so that the demand function is not partially defined in terms of its own output. Also the argument of $\Delta(R)$, and the series by which regime-pressure episodes are identified for Red Line 16 and Warning 13, under criteria published before the reading. See Section 13.7.1.

Phase gates The observable thresholds that make the adoption curve auditable. Each phase has *entry gates* (for I→II: VerifyPrice for ≥ 3 canonical workloads with ≥ 6 months of history within SLO, one corridor with zero protocol-attributable refund failures, overall refund_safe $\geq 99.5\%$ and success $\geq 95\%$, VerifyReach for major regions, ≥ 1000 receipts/day; for II→III: FeeCoverage $\geq 30\%$ of the security budget sustained 12–24 months, budgeted workloads exceeding speculative ones, VerifySettle surviving a policy shock; for III→IV: collateral use in ≥ 3 independent venues, SLOs holding through macro stress, FeeCoverage $\geq 50\%$) and *failure gates* that regress to the prior phase. Calendar horizons are intentionally omitted. The 30% gate is a phase-regression gate, distinct from Red Line 5's 10%. See Section 29.3.

Capital cycle (installation/deployment) Perez's sequence for technological surges: an installation phase financed by increasingly speculative structures, a clearing Turning Point that reprices installed capacity, and a deployment phase operating on the repriced substrate. Adopted for the transfer mechanism only — the clearing event's transfer of installed capacity from weakly financed owners to stronger users at lower capital cost — as a cross-cutting clock on Section 29.2, not as a periodization or a forecast. The thesis's instruments, not the frame, govern where they conflict. See Sources [Perez 2002].

Turning Point (capital cycle) The clearing event that ends an installation phase: financing structures written against utilization assumptions that did not arrive are repriced, and the installed base changes hands. Not a prediction of this thesis and not timed by it. Its relevance is that the fee base measured by Phase III fee coverage is plausibly a post-Turning-Point quantity (Section 29.3), and that Condition B of Red Line 14 is the trip wire for the deployment phase failing to materialize. See Sources [Perez 2002].

PIDL (PIDL) Proof Interface Definition Language — the minimal receipt schema for proofs and settlements.

PAL (PaL) Proofs-as-a-Library — SDK that compiles claims to proofs.

PRK (PRK) Privacy Rails Kit — executes non-custodial, refund-safe settlement over privacy corridors.

FER Facility Energy Receipt — signed summary of a site's energy use over a time interval.

FCR Facility Capacity Receipt — extends FER data with infrastructure resilience, grid, cooling, jurisdictional, and hardware-diversity metrics. See Section 17.7.

Delivered Verified Capacity (DVC) Scenario-specific maximum surviving flow from energy

and fuel through firm electricity, transformer and switchgear, cooling and water, hardware, network, workload, proof, verification, settlement, and usable service. Reports the active binding constraints, substitution latency, and common-cause dependencies. See Section 17.9.

Physical VerifyPrice The cost of auditing the physical infrastructure claims (FCR) behind a unit of verified work: time, cost, challenge-success rate, and open discrepancies (Section 22.8, Appendix A.8). Read by Red Line 10.

Real-Resource VerifyPrice SLO The constitutional cost band on verifying a workload's receipts, measured in real resources on reference hardware and exogenous to token price (Section 22.7). Read by Red Line 1. Distinct from Physical VerifyPrice: a network can verify proofs cheaply while its facilities are opaque, and the converse.

Sovereign Optionality ($\mathcal{O}_s$) A capacity-weighted composite of substitutability attributes—fuel diversity, firmness, interconnection, jurisdictional dispersion, and hardware mix—computed from Facility Capacity Receipt fields. It *correlates with* the number of independent physical pathways by which verified work can continue under disruption, but does not literally count them; the pathway-count question is answered by Delivered Verified Capacity. The network-level figure $\mathcal{O}_s^{\text{net}}$ is the capacity-weighted mean across facilities less a concentration penalty (an average of two capacity-share Herfindahl–Hirschman indices, over jurisdictions and over balancing authorities). Feeds risk haircuts and Red Line 13. See Section 17.8.

Topological Scarcity Lemma A nominal stock does not secure a monetary service; the relevant quantity is stress-deliverable flow through the narrowest non-substitutable edge connecting input to holder utility. See Section 8.2.

Pressure–Capacity Corridor The relation $\Pi_{\text{monetary}}(R) \propto \Delta(R) \times A_{\text{full-stack}}(R) \times H(R)$. Its likely shape is an inverted U because differential demand rises as substitutes weaken while stack availability and holder agency can fall under extreme pressure. See Section 8.2.

Disruption-Adjusted VerifyPrice The probability-weighted cost of verification across physical disruption states, rather than the observed cost under benign conditions. Used to price whether a resilience premium is buying anything. See Section 17.10.2.

Energy Interdiction The adversary class in which curtailment, rationing, tariff discrimination, interconnection denial, or long-lead equipment scarcity is used to raise the physical cost of verification without any prohibition on cryptography. See Section 5.5.

PoUW Proof of Useful Work — consensus mechanism where block rewards are earned by producing verifiable receipts of useful compute.

SLO Service Level Objective — published targets for system performance and availability.

Bronze/Silver/Gold The *latency* tiers for canonical workloads — prove-side p95 to attested result of 60 s / 10 s / 2 s at fee multipliers $1.0\times$ / $1.3\times$ / $1.6\times$ (Section 22.3.2). Assurance is a separate axis indexed by verification modality (Replicated, Replicated-plus-audit, Probabilistic, Succinct), and a quote names one of each. The same three words name the power-firmness classes of Appendix D (interruptible / curtailable-with-notice / firm), which map onto the la-

tency tiers but are not the same object; and "Gold assurance" in Section 17.5.4 is the Layer 0 sampling-evidence tier, unrelated to either.

MatMul-PoUW Proof of Useful Work construction based on matrix multiplication verification.

ZK Money An analytical lens or legacy label for service instruments referencing Privacy + Proofs. Those claims are not presumed money; only the separate base asset is a conditional monetary candidate.

Attestation Money An analytical lens or legacy label for proof-capacity and provenance-attestation demand. Formerly "Truth Money" in earlier drafts; it does not promote proof claims into money.

AI Money An analytical lens or legacy label for service instruments referencing Compute + Proofs (verified FLOPs, inference capacity). Work Credits beneath the lens remain typed service claims.

Layer 0 Verifiable Machines & Energy — open hardware and sampled supply chains as base reality.

Layer 1 Reachability — communications and transport resilience.

Layer 2 Distribution & Execution — software supply and runtime.

Layer 3 Identity & Claims — pseudonymous credentials without doxxing.

Layer 4 Truth & Work — proof systems, PoUW, VerifyPrice.

Layer 5 Value & Settlement — privacy rails and non-custodial flow.

Layer 6 Governance & Telemetry — keeping neutrality and resilience measurable.

Market Realization Plane The external financial machinery—exchanges, custody, ETFs, treasury vehicles, index products, derivatives, dealers, passive mandates, systematic and agentic allocation rules—through which claims on the native monetary object are represented and priced. Orthogonal to the stack, not a layer within it. See Section 13.7.

Native instrument A protocol-internal asset, typed service claim, evidence artifact, staking or LP position, corridor claim, or capacity voucher. Touching the protocol does not make the instrument monetary.

External financial wrapper A conventional market product written on the asset (spot ETF, ETP, treasury company, custodial balance, future, option, swap, leveraged or inverse ETP). Holding it exercises nothing. Unqualified "wrapper" means this.

Value capture vs. price capture Value capture routes native demand through the monetary object via fees, burns, and collateral. Price capture is demand for exposure to the asset's price, which can occur with no protocol use at all. Neither implies the other. See Corollary 13.1.

VerifyFlow The fourth verification family: telemetry for the external financial representation and price-transmission state of the native asset. See Appendix A.9, Section 26.8.

Flow elasticity (ε) The responsiveness of a wrapper's shares outstanding to changes in its value per share. $\varepsilon \approx 0$ means holders sit still; $\varepsilon \approx -1$ means they trim to a constant dollar

position. Measured, not structural. See Appendix H.1.

Net mechanical gain (κ) $L(L-1)+\varepsilon L^2$; the coefficient converting an underlying return into a wrapper's total return-coupled exposure demand. See Appendix H.3.

Wrapper Recycling Ratio (WRR) $-\varepsilon L/(L-1)$. Equal to 1 when holder flows fully offset gross rebalancing; below 1 the wrapper amplifies; above 1 it is countercyclical. Domain $L \notin \{0,1\}$, inverse products included. Undefined at $L = 1$: an unlevered wrapper has no rebalancing trade to recycle, so the spot ETF publishes ε itself (for which $\kappa = \varepsilon$) and the board never carries a WRR column for it. See Appendix H.5.

Wrapper Dominance Ratio (WDR) A reported *pair*, never divided into one scalar: the **stock form**, $\text{CustodialExposureShare}_t = E_t^{\text{wrapper}}/(E_t^{\text{wrapper}} + E_t^{\text{native}})$, bounded in $[0,1]$; and the **flow form**, $\text{WrapperActivityRatio}_{t,h} = V_{t-h:t}^{\text{wrapper}}/V_{t-h:t}^{\text{native}}$ over a published horizon h, unbounded above. Exposure is measured *at* an instant and usage *over* an interval, so a single ratio would scale with the window. A rising stock share with a rising activity ratio indicates the asset may be financializing faster than it is becoming money; Red Line 9 requires *both* components to rise, and publishes the split when they diverge. See Section 13.6.

Mechanical Pressure Ratio (MPR) Published in two forms that are never averaged, ranked, or combined. Gross, $\text{MPR}^{\text{gross}} = \sum_j |Q_j^{\text{mkt}}|/V_X$, counts every dollar of liquidity consumed on both sides and measures *pressure*; net, $\text{MPR}^{\text{net}} = |Q^{\text{mkt}}|/V_X$, nets offsetting flows first and measures *direction*. V_X is the asset's executable liquidity over the same window; Q^{mkt} is realized market demand, the dealer-intermediated image of exposure demand $Q^{\text{exp}} = Q^{RC} + Q^{RD}$. Read by Market Realization Warning 3. See Section 26.8.2.

Custodial Control Ratio (CCR) Base asset controlled by custodians and wrappers, divided by circulating supply. A stock, not a flow instrument; one of the two exposure components Red Line 9 reads in its full form. See Section 26.8.2.

Synthetic Exposure Ratio (SER) Delta-adjusted derivative exposure divided by liquid free float. See Section 26.8.2.

Recursive Claim Ratio (RCR) Exposure held in wrappers whose underlying is itself another wrapper or derivative, divided by total wrapper exposure, delta-adjusted on both sides over the same window as SER. One minus RCR is the share of wrapper exposure that bottoms out in the native asset directly. See Section 26.8.2.

Native Use Share (NUS) The share of observable economic activity touching the asset that actually settles through the protocol-native monetary loop: numerator, the USD value of native-loop settlement (fee-paying, proof-carrying, or corridor) over the window; denominator, all activity on public venues and wrappers reporting the asset over the same window. The universe is pre-declared with the metric; activity invisible to public data is excluded from both sides and acknowledged as a coverage limit, not assumed zero. See Section 26.8.2.

Return-coupled vs. return-decoupled flow Return-coupled flow (Q^{RC}) responds to the return itself and shapes volatility and persistence. Return-decoupled flow (Q^{RD}) arrives largely in-

dependent of the day's return and shapes destination and level. See Appendix H.7.

Wrapper–Native Growth Gap (WNG) $\Delta \ln E^{\text{wrapper}} - \Delta \ln U^{\text{native}}$, computed on unit-denominated quantities over a trailing four-quarter window (asset-price revaluation stripped from both sides), with the USD-denominated series published alongside. Persistently positive means financial exposure is outgrowing monetary use.

Red Line A protocol failure condition that falsifies the store-of-value claim if it persists without credible remediation: verification breaks, settlement safety breaks, native value capture fails, issuance becomes discretionary, censorship routes dominate, native use can be bypassed, or (Red Lines 16–18) the link-9 monetary mechanisms fail to read. Eighteen are named, each with a severity label — *kills* (the chain's hinge is severed) or *weakens* (one leg of support is removed and the thesis survives only in restated form). See Section 30.6.

Market Realization Warning A second, weaker class of signal generated by the Market Realization Plane and deliberately not a red line: it means *price has become an unreliable signal*, and it suspends the right to cite price as evidence about the protocol without necessarily saying anything about protocol function. A wrapper unwind can destroy price while damaging nothing; treating that as falsification would be as sloppy as treating a rally as confirmation. Triggers include a break in flow elasticity, a sign change in κ, MPR exceeding market-depth bounds, wrapper creations dominating native spot demand, and dealer swap capacity binding. Thirteen are named. See Section 30.7.

Service-Good Realization The pattern in which triad capacity is consumed at scale while the asset shows no insurance signature across repeated regime-pressure episodes and the monetary bid accrues elsewhere. Prescribes reclassification to *verified-capacity service asset*, not retirement. Filed as Market Realization Warning 13 because its corroborating clause reads on price; Red Line 15 separately covers the native collateral–capacity spiral. See Section 30.7.

Participation Line The household or organizational threshold below which a person or firm lacks the redundancy to act freely across time—to fail, retry, transact, move, learn, refuse coercive terms, or survive shocks. Below it, legal rights remain on paper while agency disappears in practice, and financial repression becomes politically durable. Approached along the exhaustion gradient. See Section 3.6.

Homestead Ratio The share of verified compute, proof generation, and AI-service capacity supplied by open-admission, non-hyperscaler, geographically diverse, independently verifiable operators. Its complement is read by the Enclosure Risk Flag and Red Line 11. See Section 5.6.

Agency preservation (ninth SoV requirement) The system should equip users to act privately, verifiably, and non-custodially without becoming a surveillance or dependency layer. Privacy, identity, reputation, and compliance features must be designed as tools in the user's hands: selective disclosure rather than global inspection, receipt-based reputation rather than biographies, and non-custodial settlement rather than managed accounts. A system that makes users

legible but not free has failed the political-economy test even if its cryptography is sound. Enforced on both the formal-access and the capability layer; read by Red Line 12. See Item 9.

Exhaustion gradient The approach path to the participation line: a household remains formally inside the system (banked, employed, housed, consuming) while its capacity to accumulate declines, so aggregates register nothing until a delinquency or unemployment threshold is crossed. Dollar-weighted measures lag person-weighted depletion by construction. See Section 3.6.

Household transmission The layer converting aggregate growth into household income, security, shelter, and savings. Its characteristic failure is depletion beneath resilient aggregates: consumption sustained by drawdown rather than income. Formally, operating surplus $O_i = Y_i + I_i + T_i - E_i - DS_i < 0$ financed by depleting resilience stock R_i. See Section 33.9.

Distributed resilience The requirement that system resilience be assessed jointly across the sovereign layer (R_{state}: energy, industry, defense, administration) and the household layer ($R_{\text{household}}$: shelter, savings, family formation, shock absorption), because a state can remain stable while financing continuity from household buffers — $R_{\text{state}} \neq R_{\text{household}}$. Fragility transferred downward returns through politics and defaults. See Section 33.9.

Representation discipline The publication rule that every aggregate over households or holders discloses its weighting rule, because dollar-weighted and person-weighted summaries of the same observations can move in opposite directions: $\bar{x}_{\text{dollar}} - \bar{x}_{\text{person}} = \text{Cov}(C_i, x_i)/\bar{C}$. A verified aggregate can be accurate and describe almost no participant; the weights are the explanation, and they are published. See Section 26.8.8.

Capability vs. formal access Formal access is a valid balance and the legal right to transact; capability is command over the necessities of an independent life. A ledger can supply the first while the second fails, because making claims harder to alter does not broaden who can acquire them. The ninth SoV requirement is enforced on both layers. See Item 9, Section 34.3.6.

Capital Survival Ratio (CSR) Whether successive investor cohorts in a wrapper preserved capital, as distinct from whether the product preserved AUM. See Section 26.8.4.

Policy Concentration Ratio The share of net asset demand governed by the largest common rebalancing templates, treasury algorithms, or agent policies. Ownership can be decentralized while behavior is not.

Delegation hop A transfer of a task specification across a principal–agent boundary, returning an artifact; lossy in intent in both directions (Section 7.2).

Trust perimeter The set of processes a party can sanction by means outside the protocol. A hop crosses a perimeter when the child's operator is reachable only through the protocol.

Cross-Perimeter Share (CPS) The fraction of delegation hops in a measured population that cross a trust perimeter; the demand variable for portable receipts. Not computable until receipts carry a field committing to the operator boundaries crossed.

Standing The locus of ownership, consent, liability, and revocation over an acting process;

enforceable through the stack when a person holds the key root, the delegation ceiling cannot be raised by the agent, and revocation is a protocol operation (Section 7.6).

Principal Map An overlay on the native-demand and native-holder boards that classifies each flow and cohort by ultimate principal—human with standing, firm, agent policy, or unknown—and publishes the unknown residue rather than reassigning it (Section 26.8.9).

Human-Revocable Share (HRS) The fraction of agent-controlled native balances whose key root is held by a human with standing and whose revocation is a protocol operation.

Native Policy Concentration Ratio (NPCR) The Policy Concentration Ratio computed over native positions (collateral, liquidity, reserves) rather than wrapper flows; published with exit correlation.

Standing inversion The failure kind in which agents acquire de facto standing while human title over them ceases to bind (Section 5.2).

Compositional adversary A failure mode with no malicious actor: locally rational agents whose incompatible operating rules combine into concentration, procyclicality, recursive leverage, and uninformative prices. See Section 5.2.

Collateral Loop The reflexive circuit in which asset values support tax receipts, receipts support fiscal capacity, fiscal capacity and central-bank support sustain Treasury market functioning, and Treasury yields sustain asset values. See Section 3.2.

Collateralized Sovereign Stack A description of the incumbent macro order, in which a state's effective fiscal capacity depends on the market value of the asset complexes it regulates. A diagnosis, not an architecture; distinct from the seven layers this thesis specifies. See Section 3.2.

Sovereign Credit Safety vs. Collateral Stability Credit safety is the probability of nominal repayment; collateral stability is the reliability of the instrument's price at the horizon over which it is used as margin and reference. An obligation can have the first and lack the second. See Section 3.1.3.

Automatic vs. Marginal Buyer The automatic buyer supplies demand by formula. The marginal buyer absorbs residual duration and therefore sets the price. Quantity of bids and quality of risk absorption are different objects. The two are one chain where the automatic buyer obtains its benchmark duration synthetically and leveraged intermediaries supply the other side. See Section 3.1.3, Section 3.1.3.

Description Becomes Instruction The condition in which an index that reports the market—how much eligible debt is outstanding, at what value, with what duration—comes to direct the portfolios managed against it, so that the seller of credit determines the portfolio of the buyer of credit. The human-institutional form of a holder base that is instances of one rule. See Section 3.1.3, Section 5.2.

Duration-Absorption Capacity The capacity of unlevered, automatic, and levered absorbers together to take incremental sovereign interest-rate risk (in DV01) in a period, bounded by repo funding, dealer balance sheet, margin, and volatility. The ratio of incremental DV01 to

this capacity, not the debt ratio, is where the state model locates the binding constraint. See Section 3.1.3, Section 29.2.

DV01 The dollar change in value produced by a one-basis-point change in yield. The relevant unit for whether a buyer is warehousing interest-rate risk. See Section 3.1.3.

Duration of the Claim vs. Duration of the Project Duration of the claim is the interest-rate exposure of a promised cash flow, which a repression-resistant store of value must not be. "Duration-neutral" means no contractual nominal duration; it does not mean insensitive to discount rates, liquidity, or expected utility. Duration of the project is the years between committing present resources and receiving the output of a long-lived physical asset. See Section 3.3.

Duration Warehouse An institution capable of converting present savings into long-lived productive capacity without being forced to liquidate, refinance, or reprice at the worst moment. A diagnosis of the incumbent order, not a layer of the cypherpunk stack. See Section 3.3, Section 34.5.1.

Buyer-Quality Checklist Persistence of funding, investment horizon, match to an enduring liability, capacity to bear losses without forced liquidation, and countercyclical willingness to add risk when prices fall. Published as fields, not as a multiplicative index. See Section 26.8.7.

Native Monetary Buyer Map Telemetry separating just-in-time fee acquisition, operator sell-through and inventory, burns and net issuance, self-custodied reserve demand, wrapper demand, leverage, holding period, loss-bearing capacity, and countercyclical accumulation, with every reserve outflow classed as forced liquidation, operating turnover, or exercise (spending on the service during stress, which counts for the mechanism). See Section 26.8.8.

Access Experiment The pre-registered test of whether holding the base asset in advance beats the best feasible Bitcoin-funded strategy — acquisition on the day, provider-side conversion, reserves plus service reservations — for a given workload and stress state, on equal budgets fixed before the disruption. Failure condition: Bitcoin-funded procurement reproduces the service within the same constraints without material disadvantage. A pass is not yet a reserve (Premise 7). A Phase I deliverable, not a counted red line. See Section 13.1.

Corridor Increment θ_{16}: the adjusted holder-service reading in the pre-registered moderate-pressure band minus the reading in the benign state. Red Line 16 trips when its upper uncertainty bound lies below the pre-registered minimum economically meaningful increment. See Section 30.6.

Underwriting-Model Provenance Each Red Line 17 panel lender's declaration of whether its haircut model is independently maintained, shared, templated, data-dependent, or copied. Agreement under a shared model is not evidence of symmetric knowledge; fewer than five independent approaches reads *unidentifiable*. See Section 30.6.

Project Note Explicit duration-bearing infrastructure credit with stated maturity, covenants, default states, and loss priority. Proofs may audit it; no par, redemption, or emergency-support

promise may transfer to the base asset. See Section 34.5.2.

Sovereign Maturity Transformation Financing long-lived public assets and commitments with short-term public liabilities. Duration is transferred onto the rollover calendar, not extinguished. See Section 3.2.

National-Champion Pathway The conversion of private firms into protected quasi-public infrastructure because the state's fiscal or security position depends on them. Also *enclosure by rescue*. See Section 5.2.

Numeraire-Dependence The property that a measured price path is jointly determined by the asset and by the unit it is quoted in, so the same series can support opposite conclusions about real command over resources. See Section 13.7.2.

Open Technology vs. Open Sovereignty Open-source or open-weight software is sovereign only where users can independently obtain the power, hardware, data, communications, privacy, and settlement to run it. Openness at one layer is compatible with concentration beneath it. See Section 33.4.

Strategic Tempo The conversion of physical resilience into control over *when* a disruption becomes binding, and thence into bargaining power. See Section 33.1.

Kardashev Labs The build-and-measure institution this thesis proposes to construct and instrument the stack; it issues no monetary instrument and holds no protocol treasury. See Chapter 32.

Appendix H

Formal Model of Market Realization, Wrapper Flows, and Price Capture

This appendix develops the flow machinery behind the Market Realization Plane (Section 13.7) and VerifyFlow (Section 26.8). It is deliberately quarantined here: the main body needs the distinction between value capture and price capture, not the derivations.

Attribution. The flow machinery below—flow elasticity, holder recycling, and the return-coupled versus return-decoupled distinction—is adapted from Michael W. Green's work on how product mechanics and holder behavior set prices, in particular *A Semi-Theory of Almost Everything* (Tier1 Alpha, July 2026) [Green 2026a]. Two of the results it builds on are older and are cited to their sources: the daily-reset rebalancing identity of Appendix H.2, that a fund with leverage L must trade $L(L-1)$ times its assets times the underlying return and that this holds for inverse as well as long leverage, is Cheng and Madhavan (2009); the volatility-drag result of Appendix H.9, that a daily-reset product's compound return carries a $-L^2\sigma^2/2$ variance penalty and is path-dependent, is Avellaneda and Zhang (2010). What is drawn from Green is the flow and recycling framing that sits on top of those identities: that holder flows into and out of the product are themselves return-coupled, and that the two mechanical trades can offset. The market-impact form is the square-root "law" of execution cost from the market-microstructure literature, in which impact scales roughly with the square root of order size relative to available liquidity; it is not Gabaix and Koijen's aggregate demand-elasticity multiplier, which does not appear in this model. Gabaix and Koijen (2021) is relied on only for the broader inelastic-markets premise—that flows move prices at all, because aggregate demand for financial assets is far less elastic than textbook treatments assume—which is why a flow-based telemetry layer is worth publishing. The thesis's own contribution is the generalization from equity market structure to monetary telemetry: the claim that these quantities must be published as a distinct verification family (VerifyFlow) precisely so that price movement is not mistaken for monetary adoption. Any errors in that generalization belong to the thesis, not to its sources.

A caution before the mathematics. These equations describe *transmission and magnitude*. They do not explain why a narrative ignites a behavioral shift in the first place. Flow elasticity

in particular is measured, not explained. Treating any parameter below as a structural constant is a modeling error.

Symbol collisions. Two symbols in this appendix are also used elsewhere in the thesis with unrelated meanings, and the reader should not carry either across. ε here is flow elasticity, a dimensionless regression coefficient; in Section 22.3 ε_r is the relative acceptance tolerance of the FP32 MATMUL bound. κ here is the net mechanical gain coefficient; in Section 17.8 κ_f is facility f's share of verified capacity, the weight in the network-level Sovereign Optionality aggregate. Within this appendix both symbols carry only the market-realization meanings defined below.

H.1 Flow Elasticity

For wrapper i, estimate the responsiveness of shares outstanding to changes in value per share:

$$\Delta \ln S_{i,t} = \alpha_i + \varepsilon_{i,t} \Delta \ln \mathrm{NAV}_{i,t} + u_{i,t}$$

where $S_{i,t}$ is shares outstanding, $\mathrm{NAV}_{i,t}$ is value per share, and $\varepsilon_{i,t}$ is **flow elasticity**. Interpretation: $\varepsilon \approx 0$ means holders largely sit still; $\varepsilon \approx -1$ means holders redeem roughly enough after gains to maintain a constant dollar position.

The regression has three parts and each is allocated to a different place downstream, so the decomposition must be stated rather than assumed. The ε-driven component is return-coupled and belongs in Q^{RC} below. The intercept α_i captures baseline creation or redemption that occurs regardless of the day's return — distribution growth, model-portfolio adoption, scheduled contributions — and therefore belongs in the return-decoupled term Q^{RD}, not in Q^{RC}. The residual $u_{i,t}$ stays unexplained and is published as its own series, never absorbed into either aggregate: Q^{RD} consumes α only. Assigning α_i to Q^{RC} would overstate momentum amplification; absorbing $u_{i,t}$ into either term would launder unexplained variation into a named flow. The published triple $(\hat{\varepsilon}, \hat{\alpha}, u)$ is the complete accounting; the aggregates consume two of the three.

Identification: the regression is simultaneous, and the bias is not benign. $\Delta \ln \mathrm{NAV}_{i,t}$ is not exogenous to $\Delta \ln S_{i,t}$. Creations and redemptions force dealer hedging, hedging passes through the impact function of Appendix H.8, and impact moves the underlying — which moves the NAV being used as the regressor. The OLS estimate of ε therefore suffers simultaneity bias, and its direction is the dangerous one: flows that amplify the move induce a correlation between share changes and NAV changes that OLS books as elasticity, so naively fitted $\hat{\varepsilon}$ overstates amplification precisely when amplification is present, and the fitted values then

feed κ and WRR downstream. Three disciplines follow. First, prefer instruments for $\Delta \ln \mathrm{NAV}$ that are exogenous to wrapper flow — underlying-index constituent news, macro release days, cross-venue price dislocation — over the raw regression. Second, where instruments are unavailable, publish the OLS estimate as an upper bound on amplification rather than as a point estimate. Third, never recycle a fitted $\hat{\varepsilon}$ back into the stress harness of Appendix H.12 without reporting both the instrumented and OLS values; a harness whose elasticity was fitted on the same amplification it simulates will find it.

The second bias runs the other way, and it decides when the upper bound is safe to publish. For a continuously traded underlying, the wrapper's official NAV is computed on a stale or synchronized print, not on the price at which flows actually cleared. The regressor is then a noisy measurement of the true underlying return, and classical errors-in-variables attenuation pulls $\hat{\varepsilon}$ toward zero — understating amplification. The two biases are opposite in sign: simultaneity inflates $\hat{\varepsilon}$, staleness attenuates it, and which dominates is not a property of the estimator but of the underlying's liquidity and the wrapper's pricing convention. The upper-bound publishing rule above is therefore conditional, not absolute. A fourth discipline: publish alongside $\hat{\varepsilon}$ a staleness diagnostic — the divergence between the wrapper's pricing print and the underlying's contemporaneous venue price — and treat the OLS value as an upper bound only where that divergence is small. Where staleness is material, the OLS estimate bounds nothing in either direction, and the reading must be published as *unidentified* rather than as a bound.

Because a daily-reset leveraged wrapper's NAV return is approximately $L_i r_t$ for underlying return r_t:

$$\frac{\Delta S_{i,t}}{S_{i,t}} \approx \varepsilon_{i,t} L_i r_t$$

H.2 Leveraged Rebalancing and Flow-Induced Exposure

A daily-reset fund must trade to restore its target exposure. The gross required rebalance is

$$Q_{i,t}^{\text{rebalance}} \approx L_i(L_i - 1) A_{i,t} r_t$$

where $A_{i,t}$ is fund assets. Note that $L(L-1) > 0$ both for $L > 1$ and for negative leverage such as -2 or -3: **long and inverse leveraged funds both mechanically chase the underlying move.** The inverse product is not a stabilizer.

Separately, creations and redemptions carry exposure equal to leverage times the dollar flow:

$$Q_{i,t}^{\text{flow}} \approx \varepsilon_{i,t} L_i^2 A_{i,t} r_t$$

Combining the two gives total return-coupled mechanical exposure:

$$\boxed{Q_{i,t}^{RC} \approx \left[L_i(L_i-1) + \varepsilon_{i,t} L_i^2\right] A_{i,t} r_t}$$

H.3 Net Mechanical Gain

Define the **net mechanical gain coefficient**

$$\boxed{\kappa_{i,t} = L_i(L_i-1) + \varepsilon_{i,t} L_i^2}$$

so that $Q_{i,t}^{RC} = \kappa_{i,t} A_{i,t} r_t$. The sign and magnitude of κ determine whether a wrapper amplifies or absorbs the move it sits on.

H.4 The Recycling Boundary

A wrapper fully offsets its own gross mechanical trade when $\kappa_i = 0$, which implies a critical elasticity

$$\boxed{\varepsilon_i^* = -\frac{L_i - 1}{L_i}}$$

Daily leverage L	**Full-recycling elasticity** ε^*
$+2$	-0.500
$+3$	-0.667
-2	-1.500
-3	-1.333

Table H.1: Illustrative recycling boundaries. Holder behavior stops being a vague qualitative concern and becomes an observable regime boundary.

This is the single most useful conversion in the appendix: it turns "holder behavior matters" into a number that can be measured, plotted, and breached. By construction, a $3\times$ product is self-neutralizing at $\varepsilon^* = -0.667$. If a $3\times$ product's elasticity sits near that boundary and weakens toward -0.40, the same product becomes a material momentum amplifier, and the underlying's behavior can shift from mean-reverting toward higher volatility with greater trend persistence.

H.5 Wrapper Recycling Ratio

A more legible operator metric normalizes elasticity by its own boundary:

$$\boxed{\mathrm{WRR}_{i,t} = -\varepsilon_{i,t}\,\frac{L_i}{L_i - 1}}$$

where $\mathrm{WRR} = 1$ means holder flows fully offset gross rebalancing; $\mathrm{WRR} < 1$ means the wrapper amplifies the underlying move; and $\mathrm{WRR} > 1$ means holder flows more than offset it and become countercyclical. Those readings hold where $L(L-1) > 0$ — that is, for $L > 1$ and for every inverse product $L < 0$, which together are the daily-reset classes that exist in practice — because by the identity below the sign of κ is then the sign of $(1 - \mathrm{WRR})$. A fractional-leverage product with $0 < L < 1$ would reverse the reading and is not covered; at $L = 1$ the ratio is undefined (next paragraph). This belongs on the public Market Realization & Wrapper Board.

The ratio and the gain coefficient are one object, and stating the identity prevents misreadings of both:

$$\boxed{\kappa_{i,t} = L_i(L_i - 1)\big(1 - \mathrm{WRR}_{i,t}\big)}$$

so the unnormalized quantity and its normalized image always move together, and any dashboard showing one can be checked against the other.

The unlevered case is the dominant one, and WRR is undefined there. WRR divides by $L_i - 1$, so it does not exist for an unlevered wrapper ($L_i = 1$) — and the spot ETF, the instrument most likely to become this asset's largest wrapper, is exactly that. The degeneracy is not a defect in the metric; it is the metric correctly reporting that an unlevered wrapper has no rebalancing obligation and therefore nothing to recycle. What remains meaningful at $L_i = 1$ is the elasticity itself: holder flows into a spot wrapper scale with $\varepsilon_{i,t} A_{i,t} r_t$ through the Q^{RC} term, with no rebalancing component to offset. The publication rule is therefore: **leveraged and inverse wrappers publish κ and WRR; unlevered wrappers publish ε, for which $\kappa = \varepsilon$ (dimensionless, since $L(L-1) = 0$ and $L^2 = 1$) so that $Q^{RC} = \varepsilon A r$; and the board never carries a WRR column for $L = 1$.** A reader who wants a single cross-class reading should use $\kappa_{i,t} A_{i,t} r_t$ directly, which is defined everywhere.

H.6 Wrapper-Class Applicability

The machinery above was derived for fund-like wrappers with shares outstanding, NAV, and creation/redemption mechanics — the equity-ETF structure it was adapted from. Not every wrapper class that matters for a bearer asset has those observables, and applying an equation

to a class that lacks its inputs does not produce a number; it produces a fabricated one. The publication contract is therefore class-conditional:

Wrapper class	**Principal observables**	**Applies**	**Degenerates / substitute**
Spot ETF / fund	Shares outstanding, NAV, creations	ε, κ, Q^{RD} (α-component), CSR	WRR undefined (L=1); publish ε in its place
Leveraged / inverse ETF	Shares, NAV, daily reset, L	Full apparatus: ε, κ, WRR, recycling boundary	None — this is the native habitat
Perpetual futures	Open interest, funding rate, long/short account ratio	Funding-rate response to price (a *substitute* elasticity), open-interest flow into Q^{RD}	No shares or NAV: ε via $\Delta \ln S$ vs $\Delta \ln \mathrm{NAV}$ is undefined; the leverage flows of Appendix H.2 do not arise (no reset obligation) — but position changes remain return-coupled and belong in a Q^{RC} analog built on open interest
Dated futures / forwards	Open interest by expiry, term structure	Roll flow into Q^{RD}; basis into convenience-yield telemetry (Section 26.8.6)	Same as perps; basis, not share mechanics, is the informative quantity
OTC swaps / treasury-company exposure	Disclosures, periodic filings	Delta-adjusted exposure into Q^{RD} at disclosure frequency	Stale observables; publish with the lag and never interpolate silently
Custodial / exchange balance	On-chain custodian-tagged balances	CCR, NUS, Wrapper Dominance components	No flow elasticity at all — a stock, not a flow instrument

Table H.2: Wrapper-class applicability of the VerifyFlow apparatus. A class without the observables an equation consumes does not get a fabricated reading; it gets the substitute construction named in the last column, or nothing.

Two consequences. First, the perpetual-futures substitute matters most in practice: for a crypto-native asset, perps are plausibly the largest leveraged wrapper class, and the funding-rate-to-price response is the closest available analog to elasticity — publish it, but do not label it ε, because it is not estimated from the same regression and does not carry the same interpretation. Second, the convenience-yield instruments of Section 26.8.6 already read on the derivative complex (funding, basis); the two families meet exactly here, and a reading published by

one should reconcile with the other where both touch the same venue.

H.7 Return-Decoupled Allocation Flow

Not all flow responds to returns. A thematic or passive wrapper can receive creations largely independent of the current day's return, arriving on up days and down days alike. For underlying asset j:

$$\boxed{Q^{RD}_{j,t} = \sum_k w_{jk,t}\, F^{\alpha}_{k,t}}$$

where $F^{\alpha}_{k,t}$ is the return-decoupled component of net creation into wrapper k and $w_{jk,t}$ is that wrapper's exposure weight to j. For leveraged or derivative wrappers, $w_{jk,t}$ must be a **delta-equivalent** exposure weight rather than a portfolio accounting weight, or the allocation flow will be understated by roughly the leverage factor.

The input discipline is the decomposition rule of Appendix H.1, applied here as a boundary: $F^{\alpha}_{k,t} = \alpha_k A_{k,t}$, the intercept component of the creation regression only. The residual $u_{k,t}$ is published as its own series and enters no aggregate; the ε-driven component is *excluded* because it is already inside κ_i and therefore inside Q^{RC}. Feeding raw net creation into Q^{RD} would count the same dollars twice — once as elasticity-driven exposure and once as allocation — and would propagate the overstatement into MPR, the stress harness, and any red line that reads on them. The two equations of this appendix therefore partition demand by return-dependence: **Q^{RC} owns the ε-component; Q^{RD} owns α; the residual u is published and consumed by nothing.** Where a wrapper's creation series cannot yet be decomposed, publish raw $F_{k,t}$ *flagged as undecomposed* so the aggregation layer can exclude it from Q^{RD} rather than silently double-count it.

The two flow types do different work. Return-coupled flow governs volatility, momentum, persistence, and reversal severity. Return-decoupled flow governs *destination and level*: which assets receive the structural bid, how concentrated it becomes, and which prices are accepted without valuation-sensitive selling. For a triad asset, return-decoupled flow could arrive from a spot ETF, a crypto index product, a corporate treasury mandate, an automated wealth-allocation model, a retirement default, an agentic treasury system, or a regulated "digital hard asset" basket — all potentially enormous for price and nearly irrelevant to native triad usage.

H.8 Market Impact and Liquidity

Flow becomes price only through a liquidity-dependent impact function. The standard empirical form has impact growing roughly with the square root of trade size relative to available

liquidity:

$$\boxed{I_{j,t} = Y_j \sigma_{j,t} \operatorname{sgn}(Q_{j,t}) \sqrt{\frac{|Q_{j,t}|}{V_{j,t}}}}$$

where Y_j is an impact coefficient, $\sigma_{j,t}$ is volatility, $Q_{j,t}$ is exposure demand, and $V_{j,t}$ is executable liquidity over the relevant horizon. Splitting impact into permanent and temporary components with permanence share ρ_j:

$$\Delta \ln P_{j,t} = \theta_{j,t} + \rho_j I_{j,t} + (1 - \rho_j)\, I^{\text{temp}}_{j,t} + \eta_{j,t}$$

where $\theta_{j,t}$ is the return the asset would have earned absent wrapper flow, $\rho_j I_{j,t}$ is the share of the period's impact that persists into subsequent prices, and $I^{\text{temp}}_{j,t}$ is the transient component: the same-period impact that reverts, defined as the current period's impact $I_{j,t}$ less the reversal of the previous period's transient, $I^{\text{temp}}_{j,t} = I_{j,t} - I^{\text{temp}}_{j,t-1}$ under one-period decay, so that transient terms telescope to approximately zero over a window and only $\rho_j \sum_t I_{j,t}$ accumulates into the level. Where no decay model is estimated, collapse the two terms and publish $\rho_j I_{j,t}$ alone, with ρ_j flagged as the only fitted parameter; a temporary term whose decay is not specified is not a measurement.

Fitted permanence is case-specific. Import the structure; estimate Y_j, ρ_j, and $V_{j,t}$ asset by asset.

H.9 Volatility Drag on Leveraged Wrappers

For a daily-reset wrapper with leverage L on an underlying with arithmetic drift μ (the expected simple return per unit time, so that the underlying's own log growth is $\mu - \sigma^2/2$) and volatility σ, the continuous approximation to compound growth is (Avellaneda & Zhang, 2010 [Avellaneda & Zhang 2010], for the first two terms)

$$\boxed{g_L \approx L\mu - \frac{L^2\sigma^2}{2} - (L-1)r_f - f}$$

where r_f is the financing rate on the levered portion and f is the product's expense ratio. The first two terms are the standard continuous approximation; the last two are the carry the wrapper actually bears and are frequently omitted, which flatters leveraged products in exactly the high-rate environments where they are most costly to hold. The $L = 1$ case reads correctly and should be read carefully: the Itô term $-\sigma^2/2$ does not vanish, because it is the ordinary gap between arithmetic and geometric return that any holder of the underlying bears; what vanishes is the *leverage-excess* drag $-(L^2 - L)\sigma^2/2$, together with the financing term. That excess is the quantity attributable to the wrapper rather than to the asset, and it is what the

recycling boundary and WRR of Appendix H.4 are about; those instruments, like the rebalancing identity, apply to daily-reset products with $L \notin \{0, 1\}$, since an unlevered or non-resetting wrapper has no rebalancing obligation and no excess drag.

The variance penalty grows with the *square* of leverage. At sufficiently high volatility a leveraged wrapper destroys capital even when the underlying has a positive average return. This is the formal basis for the classification in Table 8.4: a daily-reset leveraged wrapper cannot inherit the underlying asset's store-of-value status, because its terminal value depends on the path and not merely the endpoints. It is a path-dependent trading instrument that happens to reference a monetary object.

H.10 Dealer Hedge Decomposition

Product-level exposure demand is not equal to spot orders. Decompose realized hedging by instrument — physical holdings, futures, total-return swaps, options — and track delta- and gamma-adjusted exposure per counterparty, average derivative maturity, collateral and margin requirements, and evidence of counterparty caps. A migration from swaps toward options is a candidate signal that dealer balance-sheet capacity is binding, with the attendant gamma feedback when dealers are short calls. Treat this as a hypothesis to test per asset, not a law.

H.11 Measurement Contract

VerifyFlow Measurement Contract

- Use daily or intraday shares outstanding, NAV, price, AUM, holdings, leverage, and derivative disclosures.
- Chain reverse splits and corporate actions before estimating anything.
- Estimate elasticity over multiple rolling windows, never a single fixed window.
- Use structural-break tests rather than assuming fixed holder behavior.
- Address the simultaneity of elasticity estimation explicitly: instrument $\Delta \ln \mathrm{NAV}$ or publish the OLS value as an upper bound on amplification, and only where the staleness diagnostic of Appendix H.1 shows wrapper pricing tracks the contemporaneous underlying.
- Report exposure demand separately from realized spot orders.
- Delta-adjust swaps and options.
- Publish top-N counterparty and constituent concentration.
- Separate return-coupled and return-decoupled flows.
- Compare market-exposure growth against native protocol-use growth (WNG).

- Publish confidence intervals and model residuals.
- Report failed hypotheses rather than silently dropping them.

That last clause is not decoration. A flow model that only publishes the specifications that worked is indistinguishable from a narrative, and this thesis has no standing to demand receipts from protocols while exempting its own econometrics.

H.12 Reference Stress-Harness Sketch

For each scenario in Section 26.12.1: enumerate wrappers and their $(L_i, A_i, \varepsilon_i)$; compute κ_i and WRR_i; simulate a return path; accumulate Q^{RC} and Q^{RD}; map aggregate Q through the impact function given V; propagate the resulting return back into the next period's rebalance and creation terms; and record MPR, WNG, and native-use series alongside price. The output of interest is never the simulated price. It is whether the native series moved at all.

H.13 Boundaries and Caveats

- This machinery is evidence about *market transmission*. It is not evidence that Privacy, Proofs, or Compute will earn monetary premium.
- Parameters fitted in one market (equities, semiconductors) do not transfer. Re-estimate everything.
- Unlevered wrappers are not adversaries by construction. The failure mode is the inability to distinguish accessibility, custody concentration, native use, synthetic leverage, and price-insensitive demand.
- Price is never a protocol control target. The purpose of this appendix is to explain price formation, not to defend a price.
- Market realization is orthogonal to the stack. Nothing here becomes a Layer 7.

H.14 Buyer Quality and Price Provenance

The vectors above measure *product* mechanics. They do not measure *buyer* quality. The Treasury long end makes the gap precise (Section 3.1.3): a recurring dollar inflow can absorb less duration as prices fall, because market-value weights shrink, and the residual then clears with whoever is left. The same structure can appear in a monetary wrapper whose creations continue while the holders behind those creations are default-enrolled, unlevered, and duration-unaware—or leveraged, short-horizon, and one mark from a forced sale.

VerifyFlow therefore adds a buyer-quality checklist, not a new elasticity (Section 26.8.7).

For any material source of demand, record persistence of funding, investment horizon, match to an enduring liability, capacity to bear losses without forced liquidation, and whether incremental risk absorption rises or falls after a drawdown. Do not publish a multiplicative "anchor quality" index. The dimensions are the contribution; collapsing them into a scalar would recreate the scoreboard problem this appendix exists to prevent.

Approximate new DV01 absorption by an index receiving inflow F_t is

$$A_t \approx F_t \sum_j w_{j,t}\, D_{j,t}\, 10^{-4}$$

where $w_{j,t}$ is the market-value weight of instrument j and $D_{j,t}$ its modified duration. As $w_{j,t}$ falls with price, A_t can decline even as F_t is unchanged. That identity is why "inflows remain healthy" is not evidence that the residual is shrinking.

The measurement contract of this appendix applies here without modification: publish the classification of each buyer, the fields that support it, and the cases in which the classification failed. A flow model that reports only that demand arrived, and not what kind of demand it was, has not instrumented the plane.

Appendix I

Scenario Analysis: The Collateral Loop Under Stress

Epistemic Status of This Appendix

Nothing in this appendix is load-bearing. It exists to hold material that is analytically useful and evidentially weaker than the standard the main text holds itself to, so that the two are never confused.

The scenarios below are traced causal chains, not forecasts. They are offered because a mechanism worth defending against does not need to be probable, and because the discipline of writing a scenario down is what makes it possible to say later that it did not happen. Appendix I.4 lists the claims from the same source material that did *not* clear the bar, and why.

No red line in Section 30.6 and no condition in Value Capture Lemma depends on anything in this appendix.

I.1 Correlated Loss: A Recession With Rising Long Yields

The developed-market crisis template assumes a particular sequence: equities fall, capital flees to safety, Treasury prices rise, long yields fall. The sovereign's borrowing costs decline exactly when its receipts do, which is what makes the template survivable.

Section 3.1.3 gives a reason that sequence can invert. If a meaningful share of marginal demand for the reference safe asset comes from leveraged, repo-financed intermediaries, then a volatility shock does not merely fail to attract those buyers—it converts them into sellers, because their constraint is margin rather than conviction:

equity volatility rises
→ leveraged funds de-gross
→ arbitrage positions unwind
→ cash Treasuries are sold
→ long yields rise

Running concurrently, and on a slower clock:

recession
→ receipts fall + automatic spending rises
→ issuance increases
→ term premium rises

The buyer-quality argument of Section 3.1.3 supplies a third concurrent mechanism. Automatic retirement and index flows can continue in dollars while absorbing less DV01 as long-bond prices fall, and target-date rebalancing—a relative-performance rule—contributes little when equities and bonds fall by similar amounts:

stocks fall + bonds fall
→ little target-date rebalancing
+ less automatic DV01 absorption
→ greater burden on the residual buyer

The plausible path is therefore two-phase: an initial flight to safety pushes long yields down, and the balance-sheet and fiscal effects then reverse it. The scenario is not that this must happen. It is that the conditions under which it *can* happen appear assembled, per Fed staff estimates of the scale and financing of the leveraged Treasury positions involved (Barth et al., 2025 [Barth et al. 2025]; Sources), and that a portfolio or protocol treasury which assumes the standard template is unhedged against the inversion.

What this would mean for the thesis. Section 3.1 describes a playbook running debt to repression to flight-to-neutrality. That chain is unaffected in direction. What changes is the transition: the flight stage may be preceded by an interval in which conventionally safe and conventionally risky assets fall together, which is precisely the interval in which an unlevered holder of neutral collateral is advantaged and a levered one is destroyed. This is an argument for the leverage discipline in Section 26.12.1, not for any particular asset.

I.2 The AI–Fiscal Collateral Loop

Section 3.2 argues that asset values are load-bearing for fiscal capacity. The sharpest form of that argument concerns a single sector, because concentration is what turns a sectoral repricing into a sovereign one.

The chain, traced:

1. Competitive pressure or higher financing costs slow the growth of AI-related valuations.
2. Slower appreciation impairs collateral values and refinancing, before any outright price decline. The failure point is the second derivative: a financing structure built on continuously rising collateral can break when appreciation merely decelerates.
3. Capital expenditure and equity prices fall.
4. Capital-gains and stock-compensation receipts fall with them.
5. Consumption weakens through the wealth effect; recession follows.
6. The federal deficit widens.
7. Treasury issuance increases.
8. Long-term yields rise.
9. Higher yields further impair AI financing, closing the loop.

Magnitude discipline. The originating commentary attaches large numbers to step 5 in particular. Those numbers are not used here and are listed in Appendix I.4. The architecture is the contribution, and it does not require them: the loop is a loop at any magnitude, and the question a reader should carry away is whether the sector's share of index concentration, household portfolios, stock-based compensation, and construction activity is large enough for step 4 to be fiscally visible. That is a measurable question, and Appendix I.5 says where to look.

Why a monetary thesis cares. Not because it forecasts a bust. Because the loop is the mechanism that would convert a technology repricing into pressure for the state interventions described in Appendix I.3—and those interventions, not the repricing, are what threaten an open stack.

I.3 The National-Champion Conversion Pathway

Section 5.2 introduces enclosure by rescue as an adversary class. This section sketches what the instruments would plausibly be, because an adversary described only in the abstract cannot be monitored.

A rescue of a fiscally load-bearing technology sector would not resemble a 2008 bank recapitalization. Central banks can support Treasury market functioning, dealer balance sheets,

repo, bank funding, credit conditions, and general discount rates. They cannot underwrite a specific firm's margins, a private valuation, a proprietary moat, or the commercial viability of a particular facility. Support for those has to come from the fiscal and regulatory side, and it arrives as a portfolio. Several of the instruments below are already observable in some form as of this writing—export controls and tariffs shielding domestic capacity, power and interconnection priority for designated loads, and strategic government equity or warrant positions in named firms have each been taken—so the scenario is not that the instruments exist. It is that they compose into a portfolio aimed at a fiscally load-bearing sector, which is a different claim and has not been established. What such a portfolio would plausibly contain:

- Government procurement at scale, converting a speculative revenue line into a sovereign one.
- Loan guarantees and public–private infrastructure financing.
- Tax incentives and accelerated depreciation for capacity investment.
- Power, transmission, and interconnection priority.
- Export controls and tariffs shielding domestic capacity.
- Restrictions on the deployment of foreign models in regulated or public-sector contexts.
- Strategic government equity or warrant positions.
- Regulatory protection, including compliance regimes whose fixed costs favor incumbents.

Each instrument is individually defensible. The composite is the conversion of private platforms into protected quasi-public infrastructure—the same administrative-convergence structure described in Section 5.2, arriving through rescue rather than through compliance.

The relevant danger. Not that these firms are saved. That an open stack must then compete against subsidized incumbents whose compute, power priority, and regulatory position are underwritten by the sovereign, on the terms described in Section 33.9. Capability audits will not distinguish the two; exit rights will.

I.4 Quarantined Claims

The source material for this appendix contains claims that are interesting and insufficiently supported. Listing them is not throat-clearing: an argument that adopts the useful half of a source and quietly discards the rest has not disclosed its own selection rule.

I.5 Indicators: What Would Confirm or Weaken This

Consistent with Section 26.7, a scenario that cannot be checked is not worth writing. These are the observable series. They are grouped to match the loop, and none requires privileged access.

Claim	Source	Why it is not used
A rescue of the AI sector has already been arranged	Gromen [Gromen 2025–26]	Asserted without evidence; unfalsifiable as stated
Open models were released specifically to impair a named foreign laboratory	Gromen [Gromen 2025–26]	Intent is not observable from release behavior; the competitive effect is what matters and is treated separately
AI accounts for 80–90% of recent U.S. GDP growth	Gromen [Gromen 2025–26]	No verifiable derivation; the loop in Appendix I.2 does not require it
Official gold revaluation to $20,000–$30,000 to extinguish debt	Gromen [Gromen 2025–26]	A policy speculation, not a stated or observed program
Imminent closure of the U.S. capital account	Gromen [Gromen 2025–26]	Contradicted by the stated posture examined in Section 33.3
Specific Treasury-yield "pain thresholds"	Gromen [Gromen 2025–26]	Precision unsupported by any published reaction function
A gold-based system would produce a 1–2% global risk-free rate	Gromen [Gromen 2025–26]	Insufficiently specified to evaluate
Fed cuts, a bull steepener, mortgage-convexity buying, weaker dollar, stronger gold, and an equity rally will occur in that sequence	Green [Green 2026d]	A tactical path from commentary, not a state transition the thesis can assume; Section 29.2 uses quiet inflation compensation only as a tripwire for whether market repair remains available
The dominant voice in the Treasury market is now institutions following mandates and the arbitrageurs accommodating them	Green [Green 2026g]	The essay's own proposed decomposition—net duration demand by tenor and by decision rule—has not been run, and the essay says so; Section 3.1.3 adopts the mechanism from the Board's staff papers and holds the dominance claim as a conjecture awaiting that test
Any attribution of the September 2026 policy move, of the long-end selloff that followed it, or of the level of long yields, to benchmark mechanics	Green [Green 2026g]	Not claimed by the source, which distinguishes a level effect from a frequency effect; a later price move is an observation whose cause still has to be identified, and the mechanism is used for the failure mode, not for any yield
Any named energy cutoff, maritime escalation, or infrastructure attack is a base-case forecast	[Doomberg 2026c]	The physical mechanisms are useful stress scenarios; event timing, intent, and joint probability are not established
Aggregate North American energy abundance guarantees local deliverability	[Doomberg 2026b]	Conversion, location, product, transport, and timing can still bind; DVC is designed precisely because aggregate stock is insufficient
An operator targeted sponsor interests intentionally because the effects harmed the sponsor	[Doomberg 2026d]	Harm and divergent incentives do not establish

Treasury plumbing. Hedge-fund Treasury positions and their domicile; repo haircuts and funding spreads; futures basis and swap-spread positioning; dealer inventories; Treasury market depth during equity selloffs; whether long yields reverse upward after an initial risk-off decline; buybacks, issuance maturity shifts, and changes in regulatory treatment of sovereign paper.

Duration absorption (Section 3.1.3). The same plumbing read in the unit that binds. Net Treasury issuance by tenor in DV01 rather than in dollars; the option-adjusted duration and convexity of the market-value-weighted benchmark, published as an *output* series that nobody chose; asset-manager Treasury-futures longs against leveraged-fund shorts in the CFTC Traders-in-Financial-Futures report, read as a configuration and not as a census of basis trades; the cash–futures basis and implied repo; SOFR against repo and the volume of overnight Treasury repo; and the share of cash and repo activity centrally cleared, which sets how much balance sheet each unit of intermediation consumes. These are external series the thesis reads. It publishes none of them. The configuration that says the absorber is scaling is rising DV01 supply, rising benchmark duration, rising asset-manager longs and leveraged shorts, rising repo, and orderly yields. The configuration that says the mounting is failing is the same supply with rising rate volatility, funding pressure, a widening basis, contracting leveraged shorts, and deteriorating cash-market depth.

AI financing. Hyperscaler free cash flow after capital expenditure; data-center project finance terms; vendor financing and lease obligations; private valuation growth *rates* rather than levels; credit-default swap spreads; margin loans collateralized by sector holdings; delayed rounds and down rounds.

Open-model economics. Price–performance rather than benchmark performance alone; open-weight adoption among enterprises; local inference and hardware compatibility; developer migration; whether proprietary providers are forced into persistent price cuts; restrictions on foreign model deployment.

Fiscal feedback. Non-withheld individual tax receipts; capital-gains and stock-compensation receipts; corporate receipts from technology firms; consumption sensitivity to equity declines; deficit behavior during even a modest correction; whether Treasury issuance rises while private sector borrowers seek capital simultaneously.

Industrial policy. Grid generation and transmission additions; transformer and switchgear capacity; machine-tool orders; skilled-trade wages and vacancies; semiconductor, shipbuilding, and critical-mineral subsidies; trade and investment restrictions.

Physical conversion and common cause. Refinery and product throughput rather than crude stock alone; firm-power availability; transformer and switchgear outages; cooling-water restrictions; hardware and firmware dependencies; cable and network-route failures; scenario-specific DVC, active binding constraints, and substitution latency.

Settlement architecture. Official gold accumulation; use of non-dollar currencies in commodity settlement; dollar-stablecoin adoption abroad; tokenized Treasury growth; privacy and law-enforcement provisions in new payment systems; restrictions on self-custody or cross-border movement.

Native monetary buyers. Just-in-time fee acquisition; operator sell-through and inventory; burns and net issuance; self-custodied reserve accumulation; wrapper demand; leverage; holding-period distributions; loss-bearing capacity; and countercyclical accumulation during drawdowns and regime-pressure episodes.

Falsification. The collateral-loop framing is weakened if a significant equity drawdown passes through to receipts and deficits far more weakly than Section 3.2 implies, or if a volatility shock produces the textbook flight-to-safety in long yields with no subsequent reversal. Both are observable within a single cycle. Record the prediction before the event, per Section 26.7.

Pre-registered test. One test is stated here with a magnitude and a horizon so that it can be scored rather than argued about afterwards. **Trigger:** a peak-to-trough drawdown of at least 20% in the S&P 500 beginning within 36 months of this version's date. **Prediction:** over the four fiscal quarters following the trough, U.S. federal non-withheld individual income tax receipts (the series that carries capital-gains and stock-compensation realizations) fall by at least 15% against the same four quarters of the prior year, and the federal deficit widens by at least 1.5 percentage points of GDP relative to the last pre-drawdown CBO baseline, with the loop in Appendix I.2 predicting the larger part of the receipts decline to come from technology-sector realizations. **Weakened if:** the trigger occurs and non-withheld receipts fall by less than 5% while the deficit widens by less than 0.5 pp of GDP against baseline; that outcome would mean receipts are far less collateral-sensitive than Section 3.2 implies, and the framing is restated accordingly. **Not scored if:** the trigger does not occur within the horizon, in which case the test is re-registered, not counted as a pass. The companion test for the yield-inversion scenario of Appendix I.1 is that, given the same trigger, the 10-year Treasury yield stands at least 40 basis points *above* its level at the initial risk-off low within 60 trading days of that low; a textbook flight-to-safety with no such reversal scores against it.

Appendix J

What Changed Across Releases

This appendix records what each version added, in the words used when it was added. It is release history, not argument: nothing here is required to read the thesis, and no claim depends on it. It is kept because a document that revises itself in public should say what it revised.

v1.2 additions: Layer 0 now includes **Facility Capacity Receipts** (Section 17.7) and **Physical VerifyPrice** (Section 22.8); Layer 6 telemetry now includes the **Wrapper Dominance Board** and **Agency Preservation Board** (Section 26.9). A new political-economy thread runs through Part I: the **Treasury Market as Control Panel** (Section 3.1.2), the **Participation Line** (Section 3.6), **Administrative Repression** (Section 5.2), and **AI Homestead vs. Enclosure** (Section 5.6), culminating in a ninth SoV requirement, **agency preservation** (Item 9), and the closing frame **Equip, Don't Manage** (Section 36.4).

v1.3 additions: A thread that runs *around* the stack rather than through it. The **Market Realization Plane** (Section 13.7) models how external financial wrappers—ETFs, treasury companies, custodians, derivatives, dealers, and systematic allocation rules—represent claims on the native asset, and **Corollary 13.1** establishes that price appreciation and monetary adoption imply each other in neither direction. This adds a fourth verification family, **VerifyFlow** (Section 26.8), a corresponding public board, a **compositional adversary** class that requires no hostile actor (Section 5.2), a third boundary tension between institutional holdability and native monetary function (Section 5.6), a split of allocator adoption into native (Phase II-A) and wrapper-led (Phase II-B) paths, **Market Realization Warnings** distinct from monetary red lines (Section 30.7), and the formal flow model in Appendix H. Readers who want only the conceptual payload can read Section 13.2.1 and Section 36.2 and skip the appendix entirely.

v1.4 additions: A thread that runs *beneath* the stack. Where v1.3 asked who captures the value, v1.4 asks what the substrate depends on. Layer 0 already instrumented physical fragility; it did not argue it. Three additions close that gap. **Energy & Physical Interdiction** (Section 5.5) adds an adversary class that raises the cost of verification without prohibiting anything—curtailment, tariff discrimination, interconnection denial, load prioritization—and separates the verification-side exposure (reference-hardware obtainability) from the proving-side exposure (bulk power), which behave differently and are routinely conflated. **Sovereign Optionality** (Section 17.8) aggregates existing Facility Capacity Receipt fields into a single published exposure, feeding risk haircuts so that issuance discipline extends from "we cannot see

this facility" to "we can see it clearly and it is fragile." **Red Line 13** makes the dependency falsifiable: Physical VerifyPrice SLOs are exogenous to token price, as claimed, but not to the conditions under which power and hardware are obtainable—which means the hinge is externally triggerable, and should be monitored upstream rather than discovered at the hinge. Section 17.10 resolves the resulting tension between demanding cheap verification and expensive redundancy via **disruption-adjusted VerifyPrice**, with resilience explicitly subordinate to the affordability hinge. Finally, Chapter 33 adds the case the thesis previously argued only in the abstract: a state-scale **closed sovereign stack** that validates the Layer 0 premise while inverting who ends up sovereign, plus three objections it provokes (Section 34.2.6, Section 34.3.5, Section 34.3.7).

v1.5 additions: No new claims—this version makes the existing argument findable and harder to misread. An audit found that the thesis had accumulated excellent parts that the reader was left to assemble alone, so v1.5 adds a spine rather than a thread. The seven **premises** that had lived only on the project site are now stated in the document (Section 1.2), and Section 8.1 walks the full conditional chain in one place: each of the seven links, where it is defended, and the condition that would sever it—closing with Table 8.1, which rates our own argument link by link and marks two links *medium* (six as of v3.0) rather than claiming uniform confidence. Chapter 1 opens with reading paths so an allocator, a builder, or a skeptic can skip the cultural preamble and go straight to the load-bearing material. The Abstract and Executive Memo now say plainly what was previously implicit: that the hinge is externally triggerable, and that the competitor to an open stack is a competent closed one rather than fiat. Chapter 34 is regrouped from a flat list into monetary, infrastructure, political, and market-structure blocks, so the strongest objection is no longer the twelfth item in an arbitrary sequence. Table 24.3 maps each reference application to the layers it exercises, the boards it reports to, and the red lines it would trip—which incidentally shows that shipping applications cannot catch Red Lines 4 and 8. Chapter 16 gains a physical-sovereignty question kept deliberately separate from the legibility question, because a facility can be perfectly transparent about being fragile. Section 12.4 now walks all nine SoV requirements rather than the first seven, since the two it omitted—non-bypassability and agency preservation—are the two a merely useful capacity token would fail.

v1.6 additions: A dependency the thesis had described but not named. Where v1.4 asked what the substrate depends on, v1.6 asks what the *incumbent* order depends on—and finds a reflexive circuit rather than a hierarchy. The **collateralized sovereign stack** (Section 3.2) is the diagnosis: a state's effective fiscal capacity now depends on the market value of the asset complexes it regulates, so asset prices, receipts, fiscal room, and Treasury functioning support each other recursively and fail together. Section 3.1.3 identifies who actually clears the paper—leveraged, basis-sensitive intermediaries rather than patient foreign reserve managers—and separates **sovereign credit safety** from **collateral stability**, which are routinely conflated and

behave differently under stress. Section 5.2 adds a fourth adversary class, *enclosure by rescue*: a state need not seize a firm it can simply protect, once that firm has become fiscally load-bearing. Section 13.7.2 extends the Market Realization Plane with **numeraire-dependence**, since a measured price path is jointly determined by the asset and the unit it is measured in. Section 33.3 reports that the doctrine of Chapter 33 is now arriving in the West as well, so the thesis faces two state-key stacks rather than one, and Section 33.4 separates open weights from open sovereignty—released weights confer neither independent power, hardware, data, communications, privacy, nor settlement. Material that could not be verified to the standard the body requires is quarantined in Appendix I, which is labelled non-load-bearing and lists what was excluded and why. Nothing was renumbered: Layers 0–6 are unchanged, and no red line or value-capture condition depends on any macro claim added here.

v1.7 additions — the fee-level argument and a relocated monetary claim. This entry and the one that follows it record subtractions rather than additions, and both landed after the v1.6 artifact was frozen—a reader checking the v1.6 PDF will not find them there. The Value Capture Lemma's five conditions all govern *where* fees go; none governs *how large* they can be. Section 13.3 closes that gap and does not like all of the answer. A protocol fee is a wedge on turnover rather than a claim on residual profit, so a contestable market can be taxed and operator margin compression is irrelevant to fee revenue—but the sustainable fee is bounded by the marginal buyer's willingness to pay for the protocol's differential properties, net of the fact that verified compute costs *more* to produce, and that willingness is probably thin wherever a creditworthy, suable indemnitor substitutes for a proof. Bypass is accordingly a magnitude rather than a binary, and two bypass channels were not being counted: partial bypass at the intensive margin, and a fork of the protocol with a lower fee and no burn. The consequence is that fees, burns, and collateral establish a competitively priced cash-flow claim and a balance-sheet floor—both discounted-cash-flow quantities—and not a monetary premium. Section 13.3.11 relocates the monetary claim to the state-contingency of that differential value, separating two mechanisms that had been run together: the countercyclicality of the fee stream, which is a negative-beta discount-rate effect a discounted-cash-flow valuation captures exactly and which is therefore conceded to confer no moneyness; and a **holder-side service flow**, which is not a distribution at all and so has no stream to discount. The second is where the monetary claim now lives. It is then discounted honestly: the states where the hedge is most valuable are the states where the protocol's ability to supply is most impaired. The mechanism is standard; its magnitude is unsized and is presented as unsized. Section 13.3.12 shows collateral is unit-elastic in price, bounded by an accounting identity to large-cap-equity magnitudes, and subject to wrong-way risk no red line previously covered. Two chain-strength self-ratings were reduced accordingly (Table 8.1), and **Red Line 14** makes the fee-level question falsifiable, bringing the published count to fourteen.

Also in v1.7 — the Bitcoin objection, and a correction to how this document talked about

proof-of-work. The objections chapter answered gold and did not answer Bitcoin, which for a store-of-value thesis is the more consequential omission. Section 34.1.1 now opens the monetary block and concedes more than is comfortable: Bitcoin already satisfies the requirements set out here, on a fifteen-year record, through a work function with *no external buyer* — and that absent buyer is the source of its objectivity, not an inefficiency. Proof-of-useful-work weakens it. The thesis's actual claim is narrower than earlier drafts implied: not that a better money is available, but that Bitcoin does not by itself supply portable attestation or verified compute, that these are separate goods with separate demand, and that whether they carry monetary premium or price as services is open and measurable. Passages in Parts I, II, III, V, and VI that described SHA-256 work as "waste" or "heat" have been rewritten to state the trade honestly—**monetary objectivity exchanged for capacity relevance**—and a succession framing running from Part I to the conclusion has been made non-hierarchical, since the thesis claims a different function rather than a later stage. What losing to Bitcoin would look like is specified as **Service-Good Realization** (Section 30.7), and is deliberately filed as a market realization warning rather than a fifteenth red line: its price-comparative clause cannot be a red line without breaking Corollary 13.1. The count stays at fourteen.

v1.8 additions — duration of the claim, duration of the project. A completeness amendment, not an architectural one. Layers 0–6 are not renumbered. Section 3.3 distinguishes duration of the *claim* from duration of the *project*: a repression-resistant store of value must remain duration-neutral, while civilization still requires a duration warehouse capable of holding the construction interval of plants, grids, and data centers. The two must not be the same instrument. Section 3.1.3 extends the residual-buyer account from leveraged intermediaries to automatic index flows that absorb less DV01 per dollar after a selloff. Chapter 33 names China's directed banks as a duration warehouse and the Western bill-heavy book as sovereign maturity transformation. VerifyFlow gains a buyer-quality checklist, not a scalar index (Section 26.8.7). A new objection, Section 34.5.1, concedes that neutral money does not finance the reactor and refuses to let proofs, Work Credits, or the native asset become a coupon by another name. Tactical bond-market claims are quarantined in Appendix I.4.

v1.9 additions — deliverable service, monetary risk absorption, and architectural agnosticism. The core claim is narrowed: a bearer base asset may earn monetary premium only where the full service path remains stress-deliverable, native demand cannot bypass it, a persistent self-custodied loss-bearing constituency absorbs residual risk, and project credit remains separate. The steel-man expands to ten premises and the conditional chain to nine links. Section 8.2 adds the Topological Scarcity Lemma, the inverted-U Pressure–Capacity Corridor, and the Triad Coherence Test comparing one-token, neutral-reserve-plus-credit, and shared-settlement-plus-modular-collateral designs. Section 17.9 restores conversion throughput through scenario max-flow/min-cut, substitution latency, and common-cause dependency analysis; issuance is bounded by stress-adjusted DVC rather than nameplate capacity. Sec-

tion 26.8.8 moves buyer quality into the monetary chain, while the public boards add common-cause and sponsor dependencies. Section 29.2 separates four policy states without turning them into tactical forecasts. Section 30.5 compares three collateral models, and Red Line 15 makes the native collateral–capacity spiral measurable. Section 34.5.2 specifies explicit proof-audited project credit with a constitutional prohibition on transferring par, redemption, or emergency-support promises to the base asset. Work Credits are reclassified as typed service claims; FCR, FER, and PIDL artifacts as evidence; project notes as duration-bearing credit; and LP/staking positions as derivatives or operating claims.

v2.0 — the assessment pass. No new claims: this version makes the v1.9 text agree with its own falsifiability standards. A systematic adversarial review resolved five internal contradictions. Gold-tier sampling coverage is counted in production lots at tier minimum rates, with the rule-of-three converse stated (Section 17.4). A maturity gate blocks Tier A (pristine-collateral) issuance while any PoUW verification component is Experimental or Pilot—currently binding on every PoUW workload (Section 22.10). The absolute $\varepsilon = 10^{-5}$ bound on FP32 matmul verification is replaced by a scaled relative bound with a pinned reference kernel, and "transcript determinism" is restated as pinned-kernel reproducibility (Section 22.3). The Red Line 8 emergency path is reconciled as subtractive-only, with routine use of subtractive powers tripping the red line and signer accountability specified (Section 25.5.3). Naive demand-to-price passages in Part II are explicitly marked as routing rules rather than findings, with the refutation cross-referenced where it lands (Section 13.2). Four analytical gaps are given measurement machinery: a fifth measurement family, convenience-yield telemetry, prices the holder-side service flow through base-asset lending rates, forward and perpetual basis, and wrapper basis (Section 26.8.6); the joint failure mode—closed-stack services plus wrapper price realization leaving the asset a reference price with no native loop—is named with a triple-divergence diagnostic (Section 33.9); stablecoins are engaged as the incumbent bypass channel (Section 13.5.1); and the utility-token record is cited against the thesis, with the 30% fee-coverage gate labeled an order of magnitude above the precedents (Section 34.1.5). Legal posture is narrowed: a securities-law analysis confronts the Howey structure [SEC v. Howey 1946] of the value-capture design itself (Section 27.2), and the AML table is reframed as requiring legislative change rather than as a boundary regulators must respect, with the Tornado Cash precedent added to defense-in-depth (Section 27.1). Verification expectations are aligned with engineering reality: INFER_LM_70B is re-tiered to aspirational with no Sev-1 threshold until ZKML maturity lifts the gate, and Phase I expects INFER_LM_7B (Section 29.3). The penultimate chapter is retitled "Why These Could Become Money," and the base-rate problem is stated: no asset in monetary history is known to have crossed from utility cash-flow claim to monetary premium (Chapter 35). Prior-art acknowledgments are added for Layers 0–2 (TUF [TUF 2011–]/Uptane [Kuppusamy et al. 2016], Certificate Transparency [Laurie et al. 2013], OONI [OONI 2012–], Tor/I2P [Dingledine et al. 2004], BIP-324 [Mehta et

al. 2023], BOLT12 [NoBS Bitcoin 2024b], energy resource-adequacy literature [NERC annual]) with matching Sources entries. The simultaneity problem in elasticity estimation—wrapper and index flows moving NAV—is addressed in the VerifyFlow Measurement Contract (Appendix H). No red line was relaxed, no threshold was moved, and no falsification condition was withdrawn.

v2.1 — the McLuhan pass. No new claims: this version adds a framing layer, not evidence. The thesis's arguments are placed in the lineage of media-transition theory ([McLuhan 1962]; [McLuhan 1964]; [Eisenstein 1979]; [Ong 1982]; [Anderson 1983]), with every use explicitly bounded as mechanism rather than measurement. Eight placements, none load-bearing. The Executive Memo gains a closing paragraph explaining why the thesis instruments rather than forecasts: every prior restructuring of the information substrate arrived as an externality of adoption, and the consequences were nobody's plan. Section 2.0.1 adds a prior-transitions passage to Chapter 2: manuscript trust anchored in persons, print in artifacts, broadcast in institutions, and the networked era arriving with no native anchor at all—the slot this thesis proposes to fill with proof. Chapter 9 gains an interiority paragraph: the private self as a print externality (Ong), which reframes surveillance as the withdrawal of a centuries-old subsidy rather than the violation of a natural constant. Chapter 3 opens with the topology observation: truth and value move on the same topology, and both are moving from broadcast ($1 \rightarrow N$) to network ($N \leftrightarrow N$); the state as typesetter of uniform value belongs here. The canonical-workload discussion (Section 22.3) names its precedent: typographic standardization preceded and enabled interchangeable parts (Eisenstein), with the necessity claim explicitly not adopted. The Wrapper Dominance Ratio (Section 13.6) gains a reading rule: first-generation use of a new medium wears the old medium's form, so a rising WDR measures a phase, and the failure gate exists because the frame does not promise completion. Chapter 35 tacks two sentences onto the base-rate paragraph: a printed scroll is still a scroll; the precedents shipped content without structure, which is why the base rate is empty and why the empty base rate is not dispositive—and the reading earns nothing, because the chain still has to prove itself in telemetry. The VerifyPrice SLO sidebar (Section 22.7) gains the habituation argument: if a medium restructures only those who habituate to it, verification must be cheap enough to become a habit before it can become anything else, which is why the exogeneity requirement is the precondition of the medium having any effect at all. Six Sources entries are added with scope notes stating what each is and is not cited for; the "we shape our tools" line is attributed to Culkin (1967), not McLuhan. No red line was relaxed, no threshold was moved, no falsification condition was withdrawn, and no claim of the thesis now depends on the frame.

v2.2 — the instrument-hardening pass. No new claims: this version makes the two flagship measurement families agree with the standards the rest of the thesis already holds itself to. An audit of VERIFYPRICE and VerifyFlow found the core derivations sound—every equation of

Appendix H re-derives and the recycling boundary verifies numerically—but the aggregation and specification layer behind them had ten defects, now fixed. **Aggregation:** the asset-level vector's $Q^{\text{exp}} = Q^{RC} + Q^{RD}$ double-counted the ε-driven creation flow, which sits inside κ in Q^{RC} and again inside raw net creation in Q^{RD}; Q^{RD} now consumes only the intercept component, per the decomposition rule Appendix H already stated but the vector violated, with the regression residual published as its own series and consumed by no aggregate (Section 26.8.2). The Wrapper–Native Growth Gap, which feeds the triple-divergence composite and Red Line 9, gains the aggregation specification and deflation adjustment its adjacent Red Line 14 quantity already carried; an unadjusted WNG can rise on a valuation artifact alone (Section 26.8.2). The Mechanical Pressure Ratio now publishes gross and net forms, mirroring the Wrapper Dominance Ratio's stock/flow split (Section 26.8.2). **Domain:** the Wrapper Recycling Ratio is undefined at $L = 1$—the unlevered wrapper class most likely to dominate this asset—so the κ–WRR identity is stated, the $L{=}1$ publication rule fixed (publish ε; never a WRR column), and a wrapper-class applicability map added so classes without shares and NAV (perpetual futures above all) get named substitute constructions instead of fabricated readings (Appendix H.5, Appendix H.6). **Anti-gaming:** the $r(W)$ production denominator gains the measurement contract its numerator always had—reference cost models, pinned like the reference kernel, independently estimated—because a prover inflating production cost drives the ratio down and the network looks more verification-asymmetric than it is (Appendix A). **Estimation:** the OLS-as-upper-bound rule for flow elasticity was conditional, not absolute: for a stale-priced underlying, errors-in-variables attenuation biases the estimate downward, and an attenuated OLS value bounds nothing; a staleness diagnostic is now published alongside, and where staleness is material the reading is published as unidentified (Appendix H.1). **Specification:** VerifyPrice carried three tuple definitions across four sites; Appendix A now states the canonical five-field form and demotes the others to views of it. VerifyPower took the reciprocal of a vector; it is now a scalar on a pinned reference workload. $\text{VerifyPrice}_{\text{adj}}$ is annotated as a convenience summary no red line reads on, with the per-state vector load-bearing, because an expectation averages away the tail that resilience lives in (Section 17.10.3). The AffordabilityRatio's perverse direction—improving when fees rise—is named, with denominator floors and a USD-referenced settlement variant (Section 22.7). Every fix tightens an instrument; none relaxes one. No red line was relaxed, no threshold was moved, and no falsification condition was withdrawn.

v2.3 — the coherence pass. No new claims: a systematic audit of self-coherence, local coherence, and global coherence produced 125 severity-ranked findings, all remediated. **Two mathematical errors:** the USD price model multiplied Joules directly by a USD/kWh price — a dimensional error overstating the energy term by 3.6×10^6 — now derived per run with the explicit conversion and cost percentiles computed as percentiles of per-run cost, never recomputed from percentile inputs (Section 22.6); and the WNG revaluation artifact was stated

backwards — the phantom arises on a price *rise* (wrapper exposure is a stock revalued with every tick; native use accumulates at transaction prices), now stripped by unit-denominated aggregation (Section 26.8.2). **Conventions pinned once:** Q^{RD} owns the intercept α only, with the regression residual published as its own series and consumed by no aggregate; the take-rate τ uses gross fees while FeeCoverage uses retained fees, dissolving every fee-plus-burn ratio and retiring the label; the WDR is a pair and Red Line 9 requires both components to rise; two distinct deflators — WNG unit-denomination versus Red Line 14-B's workload unit-price index — are named and kept separate. **Red lines:** RL14-B gains persistence and a named deflator; RL11/RL13 thresholds instantiated where "e.g." placeholders stood; RL15's clauses labeled conjunctive; severity labels ("kills" vs. "weakens") defined; response-protocol clocks given pinned start triggers and a no-reset rule. **Instruments:** VerifyPower carries the token price explicitly ($q_{\text{fee}} \cdot P_X/p_{50,c}$); the AffordabilityRatio's undefined "floor" replaced by three concrete guards; the Part III decision rule gates on the worst disruption state's per-state bound with the π-weighted average demoted to tie-breaker; the canonical tuple reconciled across all five sites. Roughly forty definitional repairs across Parts I–VI, the appendices, and the executive memo, including the Value Capture Lemma's failure clause (no *monetary premium*, not no value at all), the Work-Credit supply identity, the θ_t double-count, the φ_t false equivalence, and the base-rate defense restated in its honest two-move form. No red line was relaxed, no threshold was moved, and no falsification condition was withdrawn; several conditions became checkable that previously were not.

v2.4 — the household-transmission pass. New claims, all bounded: this version adds the household layer to a thesis whose adversaries had been exclusively sovereign and financial. The source is Green's "The Vibecession Was Real" [Green 2026e] (Sources), adopted for mechanism and quarantined for quantities, simulations, and psychology. **The social contract as a failing soft guarantee:** the behavioral promise that discipline converts into advancement — honored contract by contract while failing in aggregate, because the incumbent's balance-sheet wealth is the entrant's acquisition price — is added to Chapter 2 as a trust failure that survives audit. **The exhaustion gradient:** the participation line (Section 3.6) gains its approach path — formally inside the system while progressively less able to accumulate — with the observation that dollar-weighted aggregates lag person-weighted depletion until a delinquency threshold is crossed. **Capability versus formal access:** the ninth SoV requirement (Item 9) gains its second layer: a valid balance is not command over the necessities of an independent life, and a ledger can supply the first while failing the second. **Representation discipline:** every board aggregate now discloses its weighting rule, because dollar-weighted and person-weighted summaries of the same observations can move in opposite directions ($\bar{x}_{\text{dollar}} - \bar{x}_{\text{person}} = \text{Cov}(C_i, x_i)/\bar{C}$); neither is the false one, and where they diverge persistently both are published (Section 26.8.8). **Household transmission in the state model:** Section 29.2 gains cross-cutting observables — saving rate, delinquency by cohort, first-time-

buyer share, underemployment, mean-versus-median net worth, sentiment against spending — with the argument that household depletion changes how states transition, not just how they feel. **Distributed resilience:** Section 33.9 gains the formal model ($O_i < 0 \Rightarrow \Delta R_i < 0$; $R_{\text{state}} \neq R_{\text{household}}$) and the requirement that the open stack evidence distributed resilience rather than sovereign continuity alone — a stability argument, not a redistribution argument. **The household amplifier:** the joint failure mode paragraph gains the nonlinear-break mechanism — resilient aggregates as the sum of two individually fragile mechanisms, upper-cohort discretionary spending and lower-cohort necessity spending, each of which fails differently under the same asset repricing. **A new objection** (Section 34.3.6): a perfect ledger hardens an unequal starting state; the response concedes the ledger point, locates the real claim in acquisition-channel breadth rather than record neutrality, and specifies the falsifier — new-participant acquisition from labor income staying negligible while incumbent concentration rises, visible in buyer-map cohorts. Five glossary entries and the synthesis's closing edge: proof tells us whether a claim is true, representation whom it describes, capability what participants can do — a stack that verifies everything and represents almost no one is a more efficient ledger for the exhaustion dynamics it was built to escape.

v2.5 — the prose pass. No new claims: this version makes the argument read like itself. An audit against the thesis's own best prose (Part I) found the gold-standard voice allocated almost entirely to the demolitions (fee incidence, the fork, the closed-stack comparison) while the constructive chapters (Work Credits, the design space, the evaluation framework) and the instrumentation core (telemetry, red lines 1–8) were left as bullet lists and spec boxes. The capability was present throughout — the audit found Part I-grade passages inside nearly every weak section — but the document's spec-first architecture kept burying them. **The red lines get their paragraphs:** Red Lines 1–8, previously bare definitionboxes with one-line shrugs, now carry the dramatization Red Lines 9–15 already had — the touchstone for RL1, refund safety as a verdict for RL2, monoculture as a single point of political failure for RL3, telemetry capture as rigging the scoreboard for RL4, the guards-in-scrip condition promoted to trip wire for RL5, the paper-gold precedent for RL6, the shrinking anonymity set for RL7, and precedent-not-heist for RL8. **The conclusion is re-ordered and its overclaim fixed:** the "Proofs will anchor truth" formulation — explicitly disavowed in Section 34.2.3 ("proofs do not create truth; they bound disputes") — becomes "Proofs will bound what can be faked"; the seven-instrument incantation at Section 36.4's peak is replaced with three human sentences each instrument makes measurable; and "Two Loops, Not One" now precedes "Equip, Don't Manage," so the thesis ends on the owner/tenant question rather than a dashboard epigraph. **The reference applications become narratives:** the four Create/Compute/Prove/Settle/Verify recitations (Chapter 24) now open with a scene — a newsroom and a fabricated clip, a treasurer under capital controls — before the flow. **Analogies deployed throughout, all historically on-thesis:** the touchstone (verification asymmetry), Lloyd's Register and the assay punch (lot sampling),

water pressure and the trunk main (DVC and min-cut), paper gold (wrapper dominance), the fire extinguisher versus the contract for its delivery (convenience yield), SushiSwap's vampire attack (the fork's moat), a thousand cars with one GPS (policy concentration), guards paid in fresh scrip (issuance-funded security), the toll versus the detour (the capturable wedge). **Structural repairs:** the Part III closing refrain, printed twice within fifty lines (the L3 close and the Part III summary), is deduplicated and the summary given its own close; the 230-word Mechanical Pressure Ratio bullet is broken into prose; the "Cross-reference:" ledger-opener in fee routing is replaced; the hedging loop (value-capture conditions restated four times, "typed service claims" recited) is compressed to statement-once-plus-pointer. No mathematical content, definition, threshold, or cross-reference was altered; every cross-reference target and every urvives.

v2.6 — the Perez pass. No new claims, no thresholds moved, no red lines relaxed: this version adds a temporal layer to a thesis whose macro model was state-based but not yet clock-aware. The source is Perez [Perez 2002], adopted for one mechanism — the clearing event that reprices installed infrastructure — and quarantined for everything else. **The capital cycle as a cross-cutting clock:** Section 29.2 gains the observation that the four states describe the sovereign system's condition while saying nothing about *when* within an infrastructure buildout that condition is being read, and that the installation/Turning Point/deployment sequence supplies the missing axis — a clock rather than a fifth state, because it advances independently of which state the system occupies. The two frameworks are shown to make the same demand of the reader: the destination is not the next observation. **How the transition window closes:** Section 29.1 gains the argument that the bridge-then-destination picture's unstated assumption — an orderly handoff — is the historically unusual case, and that the window is more likely to close by clearing event than by accrual; the demand base the triad's fee instruments measure is disproportionately a post-repricing phenomenon, so the event arrives *before* the fee base it eventually feeds — which is why the gates are metric-anchored and Gate 6 tests through stress. **Financial deflation as a second channel:** Section 13.3.2 and Section 11.1 gain the distinction between technological deflation (gradual unit-cost decline, stable utilization) and financial repricing (step-change in the price of installed capacity, utilization resetting) — observationally distinct, same effect on the volume-over-deflation bet, and capable of moving a quarter's reading by more than a year of Moore's law. **The fee base as a deployment-phase quantity:** Section 29.3 now states explicitly that Phase III fee coverage is not a smooth path from Phase II but a quantity that plausibly becomes measurable only on the far side of a financial clearing event the thesis does not predict and cannot time — a re-timing of the gate, not a relaxation; the threshold is unchanged and the regression gate still fires on the same observable. Condition B of Red Line 14 is named as the trip wire for the opposite reading: the clearing event came and went and the fee base did not arrive. **The rescue-born pathway to State 4:** Section 33.9 gains the historically better-attested birth pathway for the closed stack

— the state not choosing coordination but inheriting it, as the residual claimant on distressed strategic assets after the clearing event — which shortens the timeline, changes the binding constraint from optimizing mandate to inherited mandate, and identifies the exact juncture at which the thesis's two adversaries become one: the market-realization plane's clearing event is the closed stack's acquisition opportunity, with the triple divergence of WNG, convenience-yield telemetry, and Homestead Ratio already specified as the readable warning. Two glossary entries and one Sources entry are added with the same scope note: mechanism only, no periodization adopted, no forecasts derived, instruments govern where they conflict.

v3.0 — the moneyness pass. The first version to change the thesis's headline claim. Prior versions located monetary premium as a *residual* — value in excess of discounted cash flows — and Section 13.3.9 correctly showed that fees, burns, and collateral do not reach it. A residual definition admits no affirmative evidence, and the thesis had accordingly retreated to a single unsized mechanism with no instrument reading on it. This version replaces the residual definition with three positive mechanisms, each carrying an instrument and a falsifier, and closes the structural gap that every one of the previous fifteen red lines read on links 2 through 8 of Section 8.1 while the monetary claim lived at link 9. **Pledgeability without due diligence** (Section 13.3.10): moneyness in the collateral literature is information-insensitivity — no-questions-asked acceptance — conventionally reached through *symmetric ignorance*, which holds only until somebody looks. Cheap universal verification reaches the same property through *symmetric knowledge*, which has no such trigger, because there is no private information left to produce. The claim is scoped to unit quality rather than to fundamentals, is stated as necessary and not sufficient against its own counterexamples (graded commodities, certified diamonds, bearer bonds, warehouse receipts), and locates the differentiator in the collateralized *position* rather than the token — which supplies, for the first time in this document, a monetary justification for Layer 0 rather than an engineering one. **Denomination** (Section 13.7.3): the strongest and stickiest monetary function, measured as free-choice denomination share over contracts the protocol did not mandate, tiered by whether the counterparty sits inside or outside the accounting perimeter, with the quote/settle matrix published as the headline denominator and bitcoin's fifteen-year record of settling without denominating stated at full strength as the standing counterexample. **Verification modalities** (Section 22.2.1): the single $r(W) \leq 0.3$ ceiling is replaced by modality bands, because a workload can satisfy that ceiling and still sit in the regime where participants sample or delegate rather than check — symmetric ignorance arrived at by an SLO the network is passing. Only succinct (M1) and algebraic or probabilistic (M2) verification carry the monetary mechanism; replicated and attested verification (M3) support service delivery, Work Credits, and SLA tiers but not the pledgeability claim, because a position whose correctness rests on a committee requires the lender to form a view about that committee. The legacy ceiling is retained and relabelled as a portfolio-level engineering floor on which no monetary claim reads. **Three new red lines, all reading on link**

9: Red Line 16 gives the existing convenience-yield telemetry a clock, a threshold, and a consequence, retiring the state-contingency claim if the series proves flat across regime states; Red Line 17 trips on cross-lender haircut dispersion failing to decline as verification coverage rises, or widening under stress while receipts remain available; Red Line 18 splits into a *weakens* condition for outsiders declining to reckon in the unit and a *kills* condition for protocol participants declining to on contracts nobody required them to write in anything, with an accounting carve-out for externally mandated quoting units without which the second condition is not defensible and reverts to *weakens*. **The high-Δ tail is disaggregated** (Section 13.3.6): the tail the fee is set against is two populations, not one — supervised counterparties who need evidence and unsuable counterparties who need absence — and each one's requirement is the other's disqualifier. The corrosion runs through the shared anonymity set rather than the rail, which is why lawful privacy resolves the technical compatibility without resolving the set composition; price discrimination between the segments is available and destroys the network good it prices; revenue is accordingly not additive across the tails, and f and Q are anticorrelated through the same distribution. The criterion is offered to Section 8.3 as a discriminator between architectures, and the instrument is departure rather than census, because a set whose composition is published is not an anonymity set. **The fork is reclassified** (Section 13.3.8) from bypass channel to alternative equilibrium in a coordination game. A fork runs the same technology, so the cost handicap vanishes and the bound becomes Δ_{coord}; the fee is subject to the minimum of that and the technical bound, which is a tightening of Section 13.3.3 rather than a relaxation. Openness is shown to cut both ways — it drives the technical differential to zero and *raises* the coordination one, since Schelling points are made of common knowledge — and the moat is stated as buyable by capital rather than guaranteed. The SushiSwap observation is relocated to its correct reference class and its limits stated: one case, selected after the fact, not a base rate. Red Line 14's reference panel gains a mandatory fork row, and acceptance breadth becomes a published series. Link 7's chain-strength rating is unchanged; the magnitude is now specified as two magnitudes, both still unestimated, each with an instrument. **Section 13.2 is rewritten** in post-retraction voice rather than annotated: its question is answered affirmatively for the first time, through four channels of which the first two are conceded to be non-monetary; the "all uses of the triad must flow through the asset" overclaim is corrected to the bypass-as-magnitude position the document already holds; net supply reduction is demoted from design property to published reading; and issuance discipline is explained as the one lemma condition doing monetary work, because an asset whose supply is discretionary is information-sensitive on supply — the same conclusion Section 13.3.11 reached from the other direction, now agreeing for a stated reason rather than by coincidence. The Value Capture Lemma is reproduced verbatim and is unchanged. Chapter 33's epistemic-status sidebar is corrected: it disclaimed any load-bearing role while the chapter was simultaneously used to support Premise 1, and the corroborative reading is now stated in both places. **No chain-**

strength rating rises. Three mechanisms and three instruments were added to link 9; none of them is evidence that link 9 holds, and a rating that rises because an argument was added is self-congratulation. No threshold was moved, no red line relaxed, and no falsification condition withdrawn. Five Sources entries and six glossary entries are added, each with the scope note the house convention requires. **Late refinement to Route A, folded into this version rather than deferred, because v3.0 had not yet been released when it was made.** The *symmetric knowledge* formulation originally read that verification is free "so everyone has already looked," which asserted a universal audit that is observed nowhere and would have been refuted on first reading. The mechanism never required it: what removes the run trigger is the absence of a *discoverable edge*, which free and public verification supplies whether or not anyone exercises it, and the discipline accordingly falls on the claimant before the claim is made — the structure by which contestability disciplines conduct without entry (Baumol, 1982). Section 13.3.10 now states the equilibrium as one in which verification is constantly available and almost never used, and both abstracts, Section 22.2.1, and the two draft mirrors were corrected to match. The Grossman–Stiglitz free-rider objection — if verification is free, nobody is paid to verify — is stated against the thesis and answered from the prover/verifier cost separation, which is also shown to be the same fact as Section 13.3.9's fee concession seen from the financing side, so the two results constrain each other rather than coexisting by accident. A second invariance is claimed and bounded: verification cost is invariant to adversarial effort as well as to notional, which answers the objection that deception scales alongside checking, and holds only over claims expressible as a computation — a class explicitly narrower than the class of claims on which trust is currently failing. The compressed cypherpunk instruction *don't trust, verify* is corrected in Chapter 1 and Section 33.2 to *verify in order to trust*, on the grounds that a thesis whose success condition is no-questions-asked acceptance cannot coherently banner permanent suspicion; what both the cryptographic and the sovereign instruction minimize is unpriced dependence, not trust. Red Line 14 gains a plain-language gloss identifying Condition A as the instrument that reads whether anybody pays for the truth, and names market indifference as the failure mode with no technical signature and the one in which the thesis is most likely to be quietly wrong. Two Sources entries are added. No threshold moved, no red line relaxed, no chain-strength rating changed. The prompt for the refinement was Green [Green 2026f], whose account of the economics of verification supplied the observation that cheap checking routinely goes unperformed and two objections this section now answers. The corrections stand on Baumol (1982) and Grossman & Stiglitz (1980) rather than on that essay: the mechanism, the monetary consequence, and the synthesis with the collateral literature are this thesis's own, and no argument here depends on a paywalled source.

v3.1 — the peer-review remediation pass. No thresholds moved upward, no red line relaxed, no rating raised; several conditions are re-specified so that they can be failed, which is a tightening. An external review of v3.0 (`REVIEW-v3.0.md` in the repository) graded the thesis on

originality, accuracy, consistency, cogency, soundness, and utility and returned forty-nine line-anchored findings; every P0–P2 item is applied here. **Structure and consistency.** A duplicated section header and leaked draft comments that compiled into two consecutive sections with identical titles in Section 13.2 are removed. "Fifteen red lines" survived in the allocator reading path, the objections chapter, and the outreach surface after v3.0 added three; every count now reads eighteen, "eleven of eighteen … the four they do not reach" reads seven and names Red Lines 16–18, and bitcoin's record is stated as seventeen years. Red Line 1's condition had drifted onto the Physical VerifyPrice series that belongs to Red Line 10; it is restored to the verification-cost series. Red Line 5 and four other passages used the fee-plus-burn quantity that v2.3 retired; every coverage instance now reads retained-fee coverage per Section 25.6.6, and the 30% adoption gate is distinguished from the 10% red line (three turnover instances survived this pass and were corrected to gross native fee turnover in v3.2). Tier C carried four incompatible rules and now has one: service-grade Work Credits at $0.3\times$, collateral-ineligible. VerifyPrice was defined three ways; the five-field tuple is canonical and the physical tuples are labelled input views with memory and bandwidth published rather than priced. Two hardcoded section numbers become cross-references. **Accuracy.** The MATMUL_4096 acceptance bound used matrix ∞-norms and admitted INT8 execution and a single output entry corrupted by two hundred times its typical magnitude, both reproduced numerically; it is replaced with the componentwise backward-error bound (Higham, 2002 [Higham 2002]) at $\varepsilon_r \approx 10^{-5}$, with an adversarial acceptance suite required at registration, and the workload is split into Freivalds (M2) and succinct (M1) SKUs in a single registry table (Section 22.3.3). Freivalds' check [Freivalds 1977] is relabelled a theorem and the soundness-per-round arithmetic corrected to $r \approx 3k/n$. Lloyd's Register, a per-vessel survey, is no longer offered as a sampling precedent; acceptance sampling (ANSI/ASQ Z1.4 [DoD 1989]) and pharmaceutical lot release are, and the detection bound gains a test-sensitivity factor σ with the sampling programme re-targeted at the defects it can see. Uniform-random-bytes transports are identified as themselves a fingerprint (Wu et al., 2023 [Wu et al. 2023]) and look-like-something fallbacks are added to Layer 1. The viewing-key table promised revocation and amount-bounding that no shielded design provides; it is replaced by a disclosure-scope model with per-row mechanisms. Refund safety is redefined as a bounded-time no-loss exit measured over protocol-attributable failures; BTC↔shielded-ZEC is stated not to exist and gated; anonymity-set metrics are per-asset. The Tornado Cash narrative is completed through *Van Loon* [Fifth Circuit 2024], the March 2025 delisting, and the August 2025 Storm conviction [SDNY 2025], and the Howey walkthrough [SEC v. Howey 1946] engages the 2025 SEC staff statements [SEC 2025a] [SEC 2025b]. Fedimint [Fedimint Project 2024–25] is labelled federated custody; Bitcoin's grid independence is narrowed to what it is. The municipal-bond paragraph that contradicted "proofs must not become the bond" is removed. Culkin's venue [Culkin 1967], the authorship of the BTC↔XMR swap protocol (Gugger, 2020 [Gugger 2020]), and C2PA's parent organisa-

tion [C2PA spec] are corrected, the FEDS Note [Barth et al. 2025] is authored, Cheng & Madhavan [Cheng & Madhavan 2009] and Avellaneda & Zhang [Avellaneda & Zhang 2010] are credited for the leverage identities, and every Sources entry carries a URL or DOI, with the nine previously undated newsletter items dated from their publications' archives and one journal attribution corrected; every entry now opens with a bracketed [Author Year] key and the body cites by that key, hyperlinked to the entry, with seven entries added where the prose named a source that had none (Nakamoto, Shannon, Chesney and Citron, the two OFAC Tornado Cash actions, FATF Recommendation 16, MiCA) and nineteen landscape references marked as not cited. **Soundness.** Section 13.3.10 is scoped against Dang, Gorton and Holmström [Dang et al. 2020]: their insensitivity is a payoff property; this section's argument is adverse selection on unit quality, Bitcoin is named as the control case whose perfectly verifiable units carry 30–50% haircuts, and the prediction is narrowed to haircut *dispersion* and post-shock convergence rather than levels. Red Line 17 is re-specified to that prediction with a lender-panel maturity gate, so that it is a test Bitcoin can pass or fail rather than one it fails by construction. Red Line 18's Condition B is re-specified on net denomination of insider obligations, and the false claim that insiders face no basis risk — and that Bitcoin cannot trip B — is withdrawn: Bitcoin sits at or near tripping it and is the standing counterexample. Red Line 16 gains a minimum of two regime transitions and a pre-registered power calculation, so that a flat series kills only when the test could have seen the effect. Maturity gates are added to the concentration lines, and every red line that reads on a series now carries a numeric trigger (Red Lines 8 and 12 remained structural conditions until v3.2 gave them thresholds; Red Line 4 is binary by nature). The question the review put hardest — why a holder hedging loss of access buys the base asset rather than Bitcoin plus Work Credits — is answered once, in Section 30.6.1, and its failure is identified as the outcome Red Lines 16 and 18 already detect. The base-rate claim of Chapter 35 is restricted to equity-like claims on a fee stream and Menger's saleability account [Menger 1892] is engaged as an argument for the commodity path rather than ignored. Issuance weights above $1.0\times$, which minted more Work Credits than delivered work, are capped; the Monetary Constitution is split into a base-asset issuance envelope and Work Credit minting rules, with the capacity modulator given a scale; the zero-DVC default is resolved in favour of explicit force-majeure declaration; Delivered Verified Capacity is stated as a generalized flow; disruption-adjusted VerifyPrice is computed over the pooled distribution rather than by averaging percentiles. Emergency delisting is advisory by default. The Value Capture Lemma is restated as an accrual result with an argument, and the Market-Price Non-Equivalence Corollary is derived from the existence of the wrapper rows. Pass (2026) is engaged on the economics of proof-of-useful-work, and the usefulness–unpredictability dilemma is stated as an open problem (Section 22.4.1). **Utility.** The institution the thesis proposes is named for the first time in the source. Chapter 32 states the investment thesis as *stack first, coin optional* — the verifiable stack and its measurement institution are underwritable public infrastructure with

buyers who are not token holders; the base-asset premium is a free option instrumented for its own retirement — and supplies dated deliverables mapped to the Minimum Viable Stack, a team shape, the structure of a planning envelope funded in deliverable-gated tranches (figures held in a separate costed plan, not in this text), an entity and licensing posture, a research agenda of the seven problems this document names as open, and kill conditions for the lab itself. The hardware appendix gains a proving-cost roadmap (Appendix E.6): the cost metric distinguished from the open-market price, the February 2026 auction reading [Shankar 2026], the per-SKU proving overhead, and the modality bet stated as a slope with an eight-quarter kill condition. The KPI appendix becomes an instrument table; the glossary gains the terms the body used without defining and is reconciled with the body on Sovereign Optionality; the SDK payroll example no longer settles in a stablecoin. Three changelog claims that described edits never made are made true by making the edits. The review's grades are recorded in the repository and not here; the thesis's job is to earn better ones on the next reading.

v3.2 — the propagation pass. No threshold moved, no red line relaxed, no rating raised, no mechanism added. A second whole-thesis review (`REVIEW-v3.1.md`) found that v3.1 had applied every remediation at the site of its finding and left the document's summaries and definitions where they were, so that the body stood at v3.1 while the Thesis in Plain Language, the Chain at a Glance, the Key Definitions, the glossary, the instrument table, and the formal-definitions appendix still read as v2.x on exactly the points v3.0 and v3.1 changed. This version propagates. **Summaries.** The Plain Language chapter and the Chain at a Glance now state the three mechanisms of link 9 — pledgeability without due diligence, denomination, and the state-contingent holder-side flow — with Red Lines 16–18 as their falsifiers; "three bridges Medium" reads six, matching Table 8.1; the chain walk's link 9 and the table's premium row name all three mechanisms rather than state-contingency alone. **Definitions.** Refund safety has one definition, in Section 23.4, and every other site — Definition 5, the instrument table, the Reader's Map, the glossary, the Part I stress test — states it in those terms rather than as a hard 1.0; the two denominators (protocol-attributable for the Layer 5 SLO at 99.9%, overall for Red Line 2 and the Phase I gate at 99.5%) are stated together once. The metric's own wording at Section 23.5, which had defined the percentage over the set it should exclude, is corrected. The glossary's entry for symmetric knowledge carried the "every participant has already produced it" formulation that the v3.0 late refinement retracted; it now states the deterrence mechanism. The glossary's take rate read "fees plus burns" against Red Line 14's gross-fees rule and is corrected; the 30% phase gate reads 12–24 months everywhere; Bronze/Silver/Gold is defined as the latency axis with the assurance axis, the power-firmness classes, and Gold assurance distinguished. The Key Definitions box no longer gives Work Credits a governance weight that Section 25.5 denies them, and the energy-anchoring sentence prices production rather than verification. VerifyReach had two incompatible definitions of succ_2 and two threshold sets; Definition 4 now defers to Section 18.5, and the Layer 1 falsifier carries its uncensored-

region denominator at every site. **Cross-Part consistency.** Section 13.3.12 stated that no red line covers the collateral–capacity loop; Red Line 15 is that loop, and the passage now says so. Section 13.3.3 argued from the v3.0 single-axis tier multipliers that v3.1 replaced and now argues from the assurance axis. The fee partition was stated four ways across Parts II–V, with a Layer 0 Assurance Fund at "2% of fee revenue" in Part III and a 10% assurance budget in Part V; the reference-design partition (70% capacity providers and delegating stakers, 20% retired, 10% assurance) is pinned once in Section 25.6 and quoted everywhere else, with the Assurance Fund as the assurance budget's largest line. Three surviving "fee-plus-burn turnover" instances read gross native fee turnover, and the v3.1 entry above is corrected on that point. Red Line 8 gains a numeric clause — a single address family above 20% of time-locked governance weight for two quarters, the threshold at which Section 25.5's per-family cap binds, and emergency-path use above two invocations in four quarters — and Red Line 12 a forced-disclosure incidence threshold and a non-custodial share floor, so that the v3.1 entry's claim that every red line carries a numeric trigger becomes true and is annotated. Ten sites that presented BTC↔ZEC as a deployed corridor after v3.1 gated it now read BTC↔XMR with shielded-ZEC gated. The trademark notice names one author. The README's red-line section, headed "Fifteen" since v2.6, lists eighteen. **The two promised tables and the instrument reconciliation.** The instrument table of Appendix B, which by its own rule treats any body metric absent from it as a defect, gains rows for the three link-9 instruments — the convenience-yield vector (Red Line 16), cross-lender haircut dispersion (Red Line 17), and free-choice denomination share with the quote/settle matrix (Red Line 18) — and for the M1+M2 share, acceptance breadth, forced-disclosure incidence, governance-weight concentration, and the Native Monetary Buyer Map; Red Line 10's row carries its numbers. The modality × grade collateral haircut table that Section 22.2.1 said a later pass owed is supplied (Table 17.3), with the modality increment exceeding the grade increment in every cell but one. The scope-drift gap of Section 13.3.10 becomes a published series, off-ledger exposure disclosure, on the Value Capture Board, with no trigger of its own. **Accuracy and naming.** The introduction's vignette now says what it shows — the stack closing — and what it does not: its batch is quoted in dollars and settled in the unit, the payment-rail configuration Red Line 18 names as the null, with the base asset present as fee medium and settlement unit only; the proof fee and corridor fee are paid in the base asset. The intro's claim that receipts survive platform stripping defers to the soft-binding account of Section 20.7. The constitutional cost band on verifying receipts is renamed the *Real-Resource VerifyPrice SLO*, so that "Physical VerifyPrice" names only the FCR audit cost of Red Line 10; the receipt schema carries an assurance tier beside the latency tier. Gold's lease rates are "ordinarily" negligible, with the early-2025 episode named; Fedimint's federation posts no bond; the stablecoin-flight and lost-bitcoin claims carry sources (Chainalysis (2024), (2020)); FCMP++ is stated to have no mainnet date set at the time of writing; the corridor time-to-finality target is set at 90 minutes, above the deployed corridor's confirmation-bound floor

rather than below it; the adoption failure-gate rows for Red Lines 1 and 3 say "falsified" as the red lines do; registry identifiers carry their suffixes; VerifyReach probes are paid from the assurance budget rather than in Work Credits; the S3 sunset is subject to the DVC bound; and twenty-two Sources entries that lacked a scope note now carry one. No threshold moved, no red line relaxed.

v3.3 — the delegation-economy pass. No threshold moved, no red line relaxed, no rating raised, no link added, no layer added. Earlier versions specified how the triad is supplied and how its value might reach a base asset, and named the buyer of verification only in passing. Chapter 7 models that buyer as a *delegation*: a task handed from one machine process to another at a speed that admits no due diligence, so that a cheaply checkable receipt is the only reliability mechanism that fits inside the hop. The chapter defines the delegation hop, the trust perimeter, and the **Cross-Perimeter Share** (the demand variable for portable receipts, published as not computable until receipts carry the field that would compute it); argues that verification cost sets whether agent delegation forms an open mesh or a set of fleets; splits the buyer's coercibility by who operates it; and observes that software agents can satisfy every measured clause of the holder anchor without having standing. The consequences are distributed: an eleventh premise (Section 1.2); the ephemeral-attester paragraph (Section 2.0.1); a fifth failure kind, **standing inversion** (Section 5.2), with a standing clause on the ninth requirement; Corollary 30.1 on Red Line 1; an agent-operation row in Table 33.1; two objections (Section 34.1.2, Section 34.3.9); and clauses on links 1, 2, and 8 that name a mechanism and leave every rating where it was. The instruments arrive as an overlay rather than a thirteenth board: the **Principal Map** (Section 26.8.9) classifies native demand and native holders by ultimate principal and publishes the Human-Revocable Share, the Native Policy Concentration Ratio, and the agent-originated share by operator class; the PIDL receipt gains an optional, backward-compatible `principal_chain` context field; Red Line 12 gains a watch clause and Red Lines 17 and 18 gain reading notes, none of which alters a condition; a *candidate* Red Line 19 on principal concentration is published with placeholder thresholds, rated *weakens*, and not counted, so the thesis still publishes eighteen red lines; and the adoption chapter gains Phase II-C, one gate row, and a second named false positive, fleet-led composability. Two physical consequences close the entry. A *joint-completion* stress scenario (Section 22.9, Section 17.12) observes that the single-verification SLO does not bound a tree of delegated verifications—at the reference failure rate roughly three 340-hop trees in ten exercise a retry or refund path—and publishes a reference tree with no SLO attached; and the procurement application (Section 24.6) is extended to the agent-to-agent delegation escrow rather than a fifth application being added, so the count of reference applications stays at four.

Stated limits of this release. The addition is flattering to the thesis, which is the reason its overclaims are refused here in advance. Agents supply no precedent: no agent economy has made anything money, and the base rate of Chapter 35 is unchanged. The Bitcoin concession stands:

a mesh of operators still has a demand curve, moved operator by operator, and an external buyer is still the trade Bitcoin declined. Link 1 has a named mechanism, not evidence; the instrument that would confirm it is specified and unbuilt. Native fees paid by one operator's fleet are exposure without agency, not monetary adoption. Cheap verification makes an open mesh possible and does not make it actual. Layer 3 is where to look for the sovereignty contest, not where premium accrues. Standing inversion is the failure the boards cannot currently see, which is a claim about instruments and not about probability. No magnitude in the chapter is projected, and nothing in it bears on price. Finally, the delegation economy is a topology of use that runs on Layers 0–6; it is refused the status of a layer for the same reason the duration warehouse and the market-realization plane were.

v3.4 — the rule-of-its-own pass. No threshold moved, no red line relaxed, no rating raised, no mechanism added to the monetary claim. This version couples two halves of the buyer-quality problem that Part I had stated separately. Section 3.1.3 had the leveraged, repo-financed residual that withdraws for reasons of margin; Section 3.1.3 had the automatic buyer that follows an index weight regardless of yield. Federal Reserve staff research [Barth et al. 2024] shows they are one chain—benchmark duration-matching, obtained synthetically in futures by funds whose cash portfolios are tilted elsewhere, with leveraged funds on the other side financed in repo—and Section 3.1.3 states it: the index does not hold the duration, it orders it; the demand that appears most stable is the demand financed most fragilely; and the two withdrawals are stages of one structure. The stressed-regime behavior is cited to the Board's March 2020 study [Kruttli et al. 2021]. **Description becomes instruction:** Green's formulation [Green 2026g] is adopted for the mechanism by which a rule that enough capital follows stops describing a market and starts being one, and is read as the human-institutional case of the concentration the threat model names for agents, which is why the Native Policy Concentration Ratio is the benchmark's native analogue. **States, not a threshold:** Section 29.2 gains a cross-cutting capacity, duration absorption, locating the binding constraint in the ratio of incremental DV01 to absorber capacity rather than in the debt ratio, with the observation that a policy move taken on a filtered yield can tighten into the mounting. **Indicators:** Appendix I.5 gains the plumbing read in DV01 and by counterparty class, with the scaling and failing configurations named; all external series, none published by the stack. **Premise 9 from the reference asset:** Section 13.7 observes that the decomposition Green asks the Board to publish for Treasuries is the object the plane specifies for the base asset, built in rather than requested. **Quarantined:** the dominance claim and any attribution of a policy move or yield level to the mechanism (Appendix I.4). Three primary sources and one commentary entry added, the commentary entry stating the citation rule for its author; three glossary entries.

Coherence. A whole-document read found sentences that had drifted from what the thesis now holds, and v3.4 brings them into line. The First Principles thesis line had claimed that the triad meets all nine requirements; it now claims that the triad is engineered to hold them

under attack, with each requirement a reading and not a promise. The Part IV opener had said that settlement makes truth into money, which the sidebar a page later denies; it now says settlement. The Introduction's falsifier sidebar had read Red Line 17 on haircut levels, which that line excludes; it now reads on how far lenders disagree. The open-hardware argument in Chapter 15 had implied that a dishonest prover can forge a sound proof. It now locates the hardware threat where it lives: in witness confidentiality, in trusted setup, and in every claim that reaches past the circuit. The Conclusion now lists its programs in the order it argues for, stack and telemetry first and the monetary experiment on top of them. Its closing epigraph now assigns the Verify* families to the stack and leaves the money to its holders. Its three final lines are stated as mandates rather than forecasts. Smaller corrections were made to the openers of Chapter 35, Chapter 33, and Chapter 34, to the close of Part V, and to the vignettes. The Treasury paragraph of the plane no longer rests on a quoted remark.

Voice. The same read found the argument sound and the book uneven, and v3.4 makes it read as one document.

- **Contributions.** Section 1.3 now names eight contributions instead of four, among them the parts of the argument that are new: deliverability as a flow, accrual separated from moneyness, price separated from adoption, and the holder anchor with the delegation economy.
- **Central Claims.** They gain a fifth, that accrual is not moneyness and price is not adoption.
- **Front matter.** The three monetary mechanisms are stated once, in the Abstract. The delegation buyer, parked since v3.3, now appears in the Abstract, the Extended Abstract, and the plain-language summary, and the Part openers carry the same thread.
- **World Forces.** The wartime peg is now told beside the Macro Playbook it illustrates. The macro diagnosis now closes on the section where the triad enters.
- **Conclusion.** It ends once rather than four times, and "owners or tenants" now closes on standing as well as on balances.
- **Version history.** Sentences in the body that reported what an earlier draft had said now state the error they corrected as a present-tense argument, and the history stays here.
- **Obedience.** Standing inversion (Section 5.2) gains a second route. Every key stays with a person, and the person never uses it, because a bad outcome after following an agent is charged to the instruction while a bad outcome after overriding it is charged to the person. The route needs no new instrument. It is read by pairing the human-revocable share with the policy concentration of the same cohort. It is stated as an objection (Section 34.3.3) and answered by default privacy and the fresh start, with the limit conceded: inside audited institutions the stack cheapens reconstruction and does nothing about how blame is allocated. Layer 6 now says who owns a threshold and that an override record judges the rule rather than the signer. World Forces no longer gives the repression mechanism a motive. The quarantine now names the long-end selloff that followed the September move.
- **Practitioner questions.** The evaluation framework and the operator checklist had tested

accrual, the stack, and wrappers, and never asked the three questions the monetary claim rests on. Both now ask them: would a lender accept the asset without investigating it, does anyone reckon in it, and is holding it worth more when substitutes fail. Both also ask whether anyone holds the asset through stress, and whether issuance is bounded by stressed delivered capacity.

- **Source concentration.** Seven essays by one commentator had come to carry 33 citations, most of them repeat pointers to ideas the thesis had already adopted and named. Each is now cited where its construct is first adopted, in the quarantine table, and here, and nowhere else: 15 citations remain, and every other passage points to the thesis's own section. The one figure taken from commentary, a leveraged product's fitted elasticity, is withdrawn, as is a quoted remark that needed the commentary to source it. No mechanism, claim, or quarantine changes.
- **Register.** Contractions are expanded outside quoted speech. The authorial "we" gives way to the thesis voice, while the civilizational "we" of the opening chapters is kept. A handful of phrases from an earlier, looser register are retired.

v3.5 — the consolidation pass. No layer added, no mechanism added to the monetary claim, no rating raised. An external whole-thesis assessment of v3.4 found the infrastructure argument strong and the monetary argument one step short, and found definitions, mechanisms, and tests that did not agree with one another. Its notes were checked against the source one by one; most held, several were narrowed, one was refuted, and the decisions taken on them are recorded below. This version repairs rather than extends.

Why hold rather than convert on the day. The assessment's central question was why holding the base asset beats holding Bitcoin and converting when the need arrives. Section 30.6.1 answered Work Credits and Bitcoin as an exit asset, and never that. Section 13.1 now opens the monetary chapter with it. A held balance removes one dependency, acquisition, and the advantage exists only in the band where conversion is unavailable, too slow, or incompatible with the transaction's privacy and control terms while native settlement and the service path still work. Holding the asset buys no capacity, price, or priority; possession-based priority is deleted from the bypass defences, and priority is sold by contract. The counterfactual is the best feasible Bitcoin-funded strategy, provider-side conversion included, which is also the bypass the value-capture conditions name. The access experiment that tests it is a Phase I deliverable and an open research problem, not a counted red line.

Red lines re-specified, before any gate opened. This release does not carry the line that no red line was relaxed: each of the three re-specifications removes at least one way its line could trip and adds others, and each box lists both. Red Line 16 now tests the corridor increment the Pressure–Capacity Corridor predicts, rather than a linear coefficient on regime pressure, which an inverted U could trip at a slope near zero while the mechanism was present. Red Line 17 keeps its relative design and adds underwriting-model provenance, with an *unidentifiable* read-

ing on a clock. Red Line 18B reads short-tenor insider obligations on a level-and-persistence rule at the same 25% floor; it previously read long-dated obligations, which the denomination theory says move last, and tripped on a share that was merely flat, so a stable share above the floor could trip it. That failure path is gone. Each box publishes what now passes, what now trips, and what is unidentifiable. All three changes are prospective: no network exists, no gate has opened, and no reading is erased.

Definitions that now agree.

- **Value capture.** The lemma's five conditions are the reference design's channels, each necessary to that route and none to value capture in general; a buyback funded by fees in another unit, pro-rata distribution, and native fees without native collateral are other routes.
- **Work Credits.** A Work Credit is a prepaid claim on reserved service owed by a named provider: commit, sell and reserve atomically, deliver and verify, extinguish and issue the receipt. Issuance had been triggered by receipt acceptance while the credit promised future service.
- **Supply.** The capacity modulator on base-asset issuance is removed; audited capacity constrains service claims, not the monetary envelope.
- **The nine properties** are the specification, not a universal law of money.
- **Smaller repairs.** Collateral is a lockup that holds while operators stay, not a valuation floor. Duration-neutral means no contractual nominal duration, not insensitivity to rates. A valuation of distributions omits the holder service; a valuation of possession need not. Delivered Verified Capacity reports binding constraints, since max-flow/min-cut fails with gains or shared edges. Table 13.2 no longer calls a cell monetary validation. Asset-quoted contracts count as unit of account whatever settles them. Haircuts are predicted tight, never low. The MatMul $1 + o(1)$ is the miner's overhead against naive multiplication. Privacy rails keep payment and proof private, and the prover's view of the witness is its own threat model. The holder map counts spending the asset on the service during stress as exercise, not exit.

Structure. The Work Credits chapter is split at Economic Linkage. Its first nine sections remain Chapter 12; everything from Economic Linkage through denomination and the market-realization plane is now Chapter 13, with every label kept. The web edition serves the new chapter at `part-ii/10b-why-hold-the-base-asset`, so no other page renumbers. A full respine around the disputed inference is deferred to v4.0.

Law. Section 27.2 is dated and transaction-specific, built on the Commission's March 2026 interpretive release and the Division's September 2026 staff FAQ on buybacks and protocol-funded burns. The chain from burn to distribution to managerial act to a stronger securities case is withdrawn as a categorical claim; governance concentration and discretionary return management remain questions to examine.

Voice and print. Sentences that still narrated the revision process are restated in the present. Authorial "we" is removed from the appendices as well as the body. The section sign had

printed as "ğ" and the lemma's negation sign as "ň", both also in the v3.4 PDF.

Declined and parked. Two of the assessment's notes are declined: letting price become evidence "through a model" would loosen a standing guardrail, and the source already says "by itself"; its replacement abstract dropped the delegation buyer. Treating repression as a specified environment rather than a diagnosis is parked with the register decision it depends on.

v3.6 — the prose and verification pass. No red line or rating changed. Two passes. The prose pass rewrites openers to lead with a claim and splits long sentences in the argument chapters — the front matter, Part II, Part VI, and the delegation chapter; definitions, red-line boxes, the lemma, the Abstract, and pull-quoted sentences were not reworded. An in-project review of developments in cryptography and proof systems since 2025 was checked against the source; each fact it brought into the text was verified at its primary source.

- **Soundness.** The stated soundness assumptions name a circuit and verifier that constrain what they claim, beside the hardness and parameter assumptions. The proof-system risk cites the 2026 Zcash Orchard counterfeiting flaw [Zcash 2026] and adds delisting by system ID and a turnstile on in-asset shielded pools.
- **Rules added.** M1 publishes its security level in provable bits. Inference proofs bind committed weights to the published model. Provenance proofs require an unrevoked certificate. Tier C attestation is not relied on against a prover with physical access [Chuang et al. 2026]. The XMR leg's confirmation depth is set above the 18-block reorganization of September 2025 [Rucknium 2025]. Outflows from each in-asset shielded pool cannot exceed the value that entered it.
- **Facts brought current.** The ZKML row gives the peer-reviewed 7B measurement [Sun et al. 2024], with millisecond verification stated for wrapped proofs. MatMul proof-of-useful-work is recorded as deployed on a public mainnet [Basu 2026]. Post-quantum risk carries NIST's transition dates [NIST 2026] and Zcash's harvest-now-decrypt-later caveat [Hopwood & Grigg 2025]. The proving-cost roadmap notes the announced pause of the Boundless proving protocol [Boundless 2026]. Bit-identical transcripts across hardware are described as not available by default rather than not achievable.
- **Record.** The Komargodski–Weinstein entry lists the current version's authors, and the Pass entry quotes the paper's consulting disclosure. The introduction's red-line summary says each line retires or downgrades the thesis at the severity its box states: Red Lines 7 and 9, and Condition A of Red Line 18, downgrade rather than retire.