Next Generation Stores of Value

Privacy, Proofs, Compute

Argument Edition

A Conditional Monetary Thesis
for Verifiable Digital Infrastructure

Jason St George

Version 4.0

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.

How to cite:

> Jason St George. *Next Generation Stores of Value: Privacy, Proofs, Compute*. Argument Edition, version 4.0. DOI 10.5281/zenodo.18902696. Licensed CC BY 4.0.

About This Edition

This edition carries the case in one sitting. It is cut from the full edition of the same version, which remains the record. Most of its text is the full edition's own; the rest condenses passages the full edition states at length. Nothing here raises a rating, adds a mechanism or moves a threshold. Where a passage is condensed, the definitions, red lines and measurement contracts it relies on are the full edition's, and the condensed text defers to them.

A cross-reference marked with a dagger, such as Chapter $1^{\dagger}$, points into the full edition, version 4.0, and uses its numbering. An unmarked reference points within this edition.

What the edition leaves out is apparatus: the layer-by-layer engineering, the appendices, the scenario catalogue and most of the measurement detail. The full edition is the place to check any claim made here.

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.

So where would a monetary premium come from? Not from fees. Fees make a good utility; they do not make money. 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 liter-

ature. 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. Both hand the keys to the state. 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 or downgrade it. This thesis does not ask to be believed. It asks to be measured. Section 9.1 states the argument as a chain of nine links and names the links it rates weakest.

The Conditional Chain at a Glance

The thesis establishes its claim through a conditional chain rather than an assertion, and a chain defended across the whole thesis is easy to lose. This page states it once; Section 9.1 walks it in full, says where each link is argued, and closes with Table 9.1, 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 6)	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, not

from fee accrual: pledgeability without due diligence, denomination, and the state-contingent holder-side flow (Section 13.2, Section 15.1, Section 15.2, Section 15.3); Red Lines 16–18 read on them. The full ratings are in Table 9.1.

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.

Short Glossary

These are the main terms the argument leans on before it stops to define them, or defines only in the full edition. Each entry is condensed; the full glossary (Appendix G†) carries every term and its notation.

Access Experiment The pre-registered test of whether holding the base asset in advance beats the best feasible Bitcoin-funded strategy, 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 7.1.

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.

Delivered Verified Capacity (DVC) How much usable service survives a given scenario, traced as a flow from energy through power, cooling, hardware, network, proof and settlement, and reported with its binding constraints, substitution latency and common-cause dependencies. See Section 17.9†.

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. Scoped throughout this thesis to *unit quality* rather than to fundamentals: an information-insensitive asset can be violently volatile. See Section 15.1.

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†.

Red Line A protocol failure condition that falsifies or weakens the store-of-value claim if it persists without credible remediation. 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 17.6.

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—conditioned on 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). The name borrows a mechanism; it does not make a triad asset a storage commodity. It is not the fee stream's countercyclicality, which a discounted-cash-flow valuation captures exactly, nor a product of burns, collateral or fee share. Unsized. See Section 15.2.

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. What it disciplines is the claimant, not the counterparty: a lie anyone could expose for free is not worth telling, so checking stays available and mostly unused, and its cost is what to measure. Its opposite, *symmetric ignorance*, holds only until somebody looks; that is the run. See Section 15.1.

Triad Privacy, Proofs, and Compute: three cryptographic service capacities that may support a monetary candidate. Whether they should share one asset is tested, not assumed.

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 (a required fee medium, 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.

VerifyPrice(W) A public KPI vector measuring the cost and time to verify workload W (Appendix A†).

Work Credit A typed, prepaid claim on reserved future service, owed by a named provider or pool and 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 10.1.1.

Contents

About This Edition ii

Abstract iii

The Conditional Chain at a Glance v

Short Glossary viii

I The Problem 1

1 Introduction 2
1.1 Central Claims . . . 6
1.2 Premises . . . 8

2 The Failure of Soft Guarantees 10
2.0.1 A Prior Transition: What the Press Actually Changed . . . 11

3 The World Forces New Monetary Primitives 13
3.1 Macro Playbook: Debt Repression Flight to Neutrality . . . 13
3.1.1 Then and Now . . . 14
3.1.2 The Treasury Market as Control Panel . . . 16
3.1.3 Who Actually Buys the Paper . . . 17
3.2 Social: The Web's Trust Default Has Flipped . . . 22
3.3 Political: The Participation Line Has Broken . . . 23

4 First Principles: What a SoV Must Survive 25

5 Threat Model 29
5.1 Five Kinds of Failure . . . 29
5.2 Declining to Energize . . . 31
5.3 Boundary Conditions . . . 31

6 The Delegation Economy: Who Buys Verification 33
6.1 The Marginal Buyer . . . 33
6.2 The Delegation Hop . . . 34

6.3 Verification Cost Sets Delegation Topology . . . 35
6.4 The Buyer Bifurcates . . . 36
6.5 Agents as Holders . . . 37
6.6 Standing . . . 38

II The Alternatives 40

7 Why Not Something Already Here? 41
7.1 Why Hold Rather Than Convert on the Day? . . . 41
7.2 Why Users Cannot Simply Bypass the Asset . . . 44
7.2.1 The Incumbent Bypass Channel Is Not Hypothetical . . . 45
7.3 Wrapper Dominance Risk . . . 46
7.4 Bitcoin, Gold, Toll Roads and the Starting State . . . 48
7.4.1 "Why not just Bitcoin?" . . . 48
7.4.2 "Gold still beats this because gold needs no network." . . . 50
7.4.3 "Why not just own the toll roads?" . . . 51
7.4.4 "A perfect ledger hardens an unequal starting state." . . . 52

8 The Closed Sovereign Stack 54
8.1 The Mirror: The Same Doctrine, Arriving in the West . . . 55
8.2 The Inversion: Who Is Sovereign . . . 57
8.3 What This Warns . . . 59

III The Instruments 60

9 The Triad and the Monetary Candidate 61
9.1 The Conditional Chain, End to End . . . 62
9.2 Topological Scarcity and the Pressure–Capacity Corridor . . . 67
9.3 Monetary Objects and Value Capture . . . 69
9.4 Three Reference Designs . . . 73

10 Work Credits: Reserved Service Claims 75
10.1 Definition: Work Receipts vs. Work Credits . . . 75
10.1.1 The Work Credit Lifecycle . . . 76
10.2 Why Work Credits Are Typed Service Claims, Not Presumed Money . . . 77

IV Delivery **79**

11 Layers of the Cypherpunk Stack **80**
11.1 Modules in Outline: Applications and Primitives 83

12 How the Stack Delivers **85**
12.1 What Each Layer Hands Up 85
12.2 Verification Asymmetry 86
12.3 VerifyPrice 87
12.4 Delivered Verified Capacity 88
12.5 Telemetry Anyone Can Re-Run 89

V Why Hold **91**

13 Why Hold the Base Asset? **92**
13.1 Economic Linkage: How Triad Demand Becomes Asset Value 92
13.1.1 Value Capture vs. Price Capture 96
13.2 Fee Incidence and the Level of the Fee 97
13.2.1 Incidence: A Fee Is a Wedge on Turnover, Not a Claim on Profit 98
13.2.2 The Base, and Why It Deflates 99
13.2.3 The Bound on the Sustainable Fee 100
13.2.4 Bypass Is a Magnitude, Not a Binary 101
13.2.5 What Determines Δ, and Why It Is Probably Thin 101
13.2.6 The Fork Is an Equilibrium, Not a Competitor 103
13.2.7 What the Five Conditions Establish, and What They Do Not 105
13.3 From Utility Demand to Monetary Premium 107

VI When It Would Be Money **108**

14 Why a Base Asset Behind These Could Become Money **109**
14.1 As a Store of Value 111
14.2 As Stack 112
14.3 As Telemetry 113
14.4 Trading One Base Reality for Another 113

15 The Three Monetary Mechanisms **114**
15.1 Pledgeability Without Due Diligence 114
15.2 Where a Premium Would Come From 119
15.3 Denomination 121

16 Stages, Phases and Macro States **124**
16.1 From Need to Price . . . 124
16.1.1 A Three-Stage Model . . . 124
16.1.2 The Negative Case: Useful Infrastructure Without Money . . . 125
16.2 Four Phases of Adoption . . . 126
16.3 Macro Policy States . . . 129

VII The Evidence **131**

17 Risk Analysis & Failure Modes **132**
17.1 Technical Risks . . . 132
17.2 Economic Risks . . . 133
17.3 Political Risks . . . 133
17.4 Spec Drift, Vaporware, and Scoreboard Capture . . . 134
17.4.1 Scoreboard Capture in the Market Realization Plane . . . 134
17.5 Collateral Architecture Under Wrong-Way Risk . . . 135
17.6 Eighteen Ways to Be Wrong . . . 135
17.7 The Three Monetary Red Lines . . . 137

18 Kardashev Labs: Build and Measure **144**
18.1 What Is Being Bought: Stack First, Coin Optional . . . 144
18.2 Research Agenda: The Problems Named as Open . . . 145
18.3 How the Lab Fails . . . 146

19 Conclusion: A Bell Labs for Privacy, Proofs, and Compute **147**
19.1 A Quiet Arrival . . . 148
19.2 Two Loops, Not One . . . 148
19.3 The Question Under Test . . . 149
19.4 Equip, Don't Manage . . . 149

Sources **151**

Part I
The Problem

The guarantees that money, media and credentialed authority ran on are failing under debt, AI and platform control. This Part says why, what a store of value would have to survive in their place, who will attack it, and who will keep buying the services that might support it: increasingly, a machine delegating to another machine.

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 full edition's front matter, then Section 1.2 (the eleven premises), Section 9.1 (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 17.6 (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 11 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 5.3 states the boundary conditions no engineering removes, Chapter 7 answers the strongest known objections (including the completeness objection in Section 34.5.1†), Chapter 8 concedes what a competent closed sovereign stack does better, and Table 9.1 rates the argument link by link.

If you want the whole claim in one sentence: *Thesis in Plain Language*, in the full edition's 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, and the attacker is no single conspirator: it is 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 a 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. For that viewer, footage that looks real is no longer evidence that anything happened.

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.

What can serve as an anchor once appearances are cheap to fake? 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.

The credo is usually compressed into a slogan, *don't trust, verify*, and the slogan 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 15.1 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 the terms: 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 inference. 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 7.4.1, and a reader who wants the concession before the argument should go there first.

What follows translates that intuition into claims, 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 17.

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 7.1).

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 applications 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 more than 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 7.4.1).

> **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, Chapter 9).

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 9.2, 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 17.6).

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 17.6).

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 6, 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.

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. The guarantees now fraying are concrete promises: that a custodian will honor a withdrawal, a platform will carry speech, a court will enforce an indemnity, a hardware vendor will keep a device usable, and a state will leave practical exit open.

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. Nobody cheated. Nobody defaulted. A trust failure that survives audit is the hardest kind to repair, and it feeds the participation line directly (Section 3.3).

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.

Orwell's boot is the repression everyone recognizes. The kind this thesis is built for arrives as paperwork: a custody default, a compliance rule, a revised term of service, each reasonable on its own page, and all of them governing the channel rather than the conclusion. Nobody is arrested. The payment just doesn't clear.

> *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 with a different ontology: what counts as real work, and how do we know?

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]. The media-transition literature [McLuhan 1964] documents the pattern as mechanism rather than measurement, and it earns 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)). Broadcast 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, which is 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 6). 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. The anti-memory-hole requirements of Chapter 26† already specify that, and this observation gives them a reason. 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. 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 shows 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

Repression rarely announces itself as a taking: it works best when holders are legible and takings are not, and the stack this thesis tests is built for the reverse. This chapter traces the pressure from the debt arithmetic that makes repression attractive, through a web whose trust default has flipped, to a participation line that has broken.

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 $\sim$2.4$\times$ multiple of world GDP [Gaspar et al. 2025]. Some economies (such as Japan) exceed 400% of GDP when public and private debt are combined. Standard scenarios from multilateral institutions point to global public debt approaching $\sim$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.

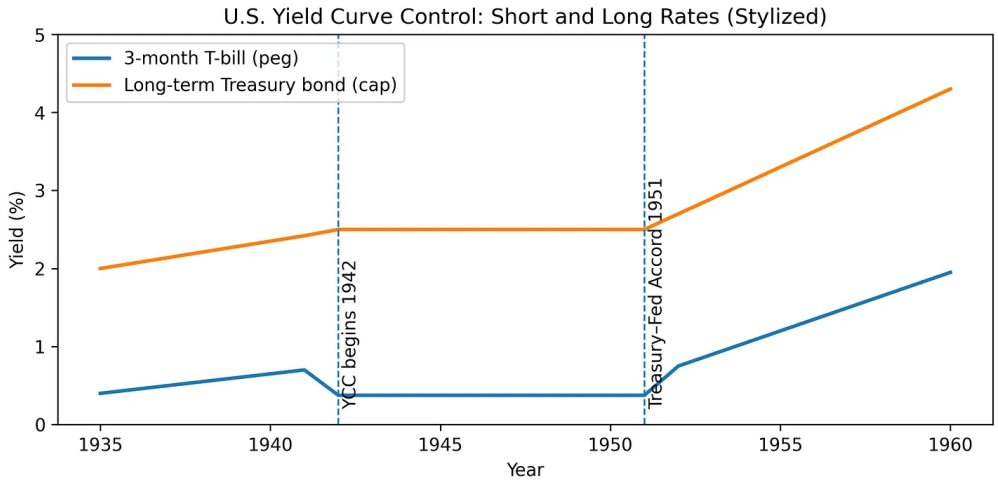


Figure 3.1: U.S. Treasury bill peg and long-bond cap under WWII-era yield-curve control (stylized). Peg at 3/8% on bills, cap at 2½% on long bonds; dashed line marks the 1951 Treasury–Fed Accord. Source: Rose (2021), Federal Reserve Bank of Chicago.

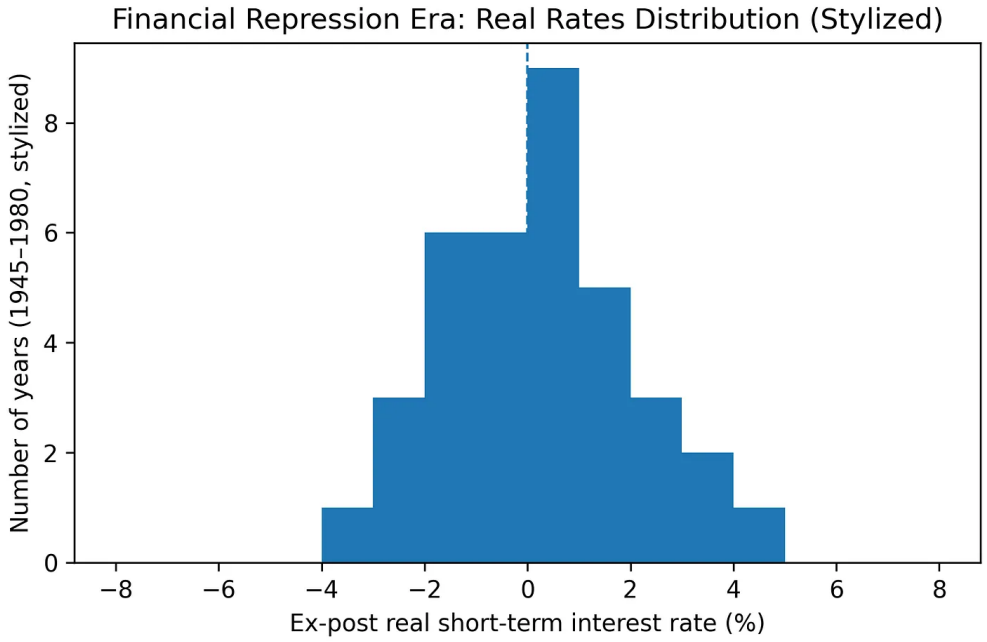


Figure 3.2: Stylized distribution of ex-post real short-term interest rates in advanced economies, 1945–1980. Real rates are negative in roughly half the years, consistent with Reinhart & Sbrancia's estimates of the financial-repression era.

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.

What the person taxed is permitted to notice. Repression's instruments, except the freeze, share one property: none announces itself as a taking. Negative real rates arrive as a price level, balance-sheet repression as a capital rule, and energy interdiction (Section 5.5†) as a queue position. They work best when holders are legible and takings are not. The stack is built for the reverse: holders private by default, takings visible. Its issuance follows a schedule fixed ex ante while Condition 4 of the Value Capture Lemma holds, so a change to it cannot be made silently. A flaw in a shielded circuit could still mint silently; the turnstile on each in-asset shielded pool caps what can leave it, though not dilution inside it (Section 17.1). The stack's telemetry publishes verification cost and corridor health (Section 22.2†, Section 23.5†), and selective-disclosure proofs serve lawful tax reporting (Section 5.4†). None of this exempts anyone from taxation or prevents coercion (Section 5.3). For holdings and services on the stack, it narrows the invisible channel and leaves the visible ones. Energy interdiction stays hard to attribute: telemetry shows the cost of proving rising, not who raised it.

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.

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 monitor not only whether users *can* access the asset, but whether they are using the monetary rail or merely holding a compliant shadow claim. Section 7.3 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 receives 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). They also find 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†). It is also 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. That intermediation is 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 16.3 specifies states with observables rather than a debt ratio with a trigger, and why the plumbing series in Appendix I.6† 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. 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.5† 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 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 a feed is simply "what's popular" no longer survives contact with the evidence. What a user sees (and does not see) is the result

of political, institutional, and algorithmic priors that user does 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 systems: 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: $O(n^2)$ to verify against $O(n^3)$ to 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, but 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.

3.3 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 6.6 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 19.4 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

What must a credible store of value survive? 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, which is to say no *contractual* nominal duration. That 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. That means a human-held key root, a delegation ceiling the agent cannot raise, and revocation as a protocol operation (Section 6.6, 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. They also 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 7.3).
- **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 three repression stress tests, written as operator SLOs that treasuries and allocators can actually underwrite and monitor:

- **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, but not before the thesis is explicit about who will attack this system and at what layers. Chapter 5 states that threat model, Chapter 11 lays out the layered architecture that the rest of the thesis will fill in, and Chapter 6 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

Every architecture is a bet about who will attack it and how. The conditions set out so far (verification cheap and public, the full path deliverable, holder agency intact under pressure) become falsifiable only once that bet is written down. The written bet names what must be protected, who is trying to break it, and which measurement decides whether the system is still credible. 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.

What has to survive is the whole Create/Compute Prove Settle Verify loop and the machines beneath it. Receipts must not be silently corrupted, flows must stay confidential past the first audit request, admission must stay open and settlement unseizable, and verification must stay cheap, because otherwise "anyone can verify" collapses into "someone we have to trust."

5.1 Five Kinds of Failure

Who attacks? Not a cartoon hacker. The primary adversary is the rational sovereign under stress. A state that must service promises larger than its productive base reaches for the levers it actually controls (rate caps, yield-curve control, capital controls, regulated custody, mandatory "secure" hardware and identity) and deputizes banks, clouds and app stores as enforcement arms. Platforms facing regulatory risk centralize moderation and provenance beside it. No one actor has to become an obvious villain; their combined incentives re-create a soft but totalizing chokepoint architecture. Beneath that axis sit economic adversaries who prefer rent to work. Beneath them, like bedrock or landmines, lies the supply chain: if the machine can corrupt inputs or falsify energy receipts, the economic claims built on it become unreliable even if the proof system remains mathematically sound. Post-quantum risk runs on a published clock [NIST 2026].

The hostile playbook is well documented, and each canonical move (filtering and active probing, app-store pressure, real-name registration, shutdowns, approved-custody defaults, stablecoin-only institutional flows) has a counter-move mapped in Table 5.1†. The pattern is constant: neutrality, privacy and verifiability are held as service-level objectives, not vibes.

Hostile capture is one of five kinds of failure, and the other four are quieter.

Institutional co-option needs no ban (Section 5.2†). Banks, platforms, custodians and regulators converge on safety, convenience and fraud prevention; each measure is defensible in isolation, and together they leave the optional rails optional only on paper. The question is not only whether the protocol can run if banned but what share of real usage stays non-custodial, private, verifiable and protocol-native when the compliant alternative is easier, which is the substitution the Wrapper Dominance Ratio measures (Section 7.3).

Mechanical fragility needs no intent at all (Section 5.2†). Sponsors maximize fees, funds follow mandates, dealers hedge and trend systems follow prices, each rational inside its own boundary, while the machine they compose produces concentration, procyclicality and false price signals. In an agent economy this is the default. A thousand human treasurers with a thousand rebalancing habits produce a distribution; a thousand agents built on one template produce a step function.

Enclosure by rescue arrives disguised as a solution, and it is the most dangerous class for this thesis (Section 5.2†). A state whose fiscal capacity rests on a private asset complex protects it when it is threatened, with procurement, guarantees, power priority and regulatory shelter. Every other adversary leaves the open stack's case intact or strengthens it; repression makes neutral rails more attractive. This one produces a competitor that is well capitalized, technically excellent, energy-privileged and genuinely useful, and it wins on capability and convenience, the ground on which the open stack concedes it cannot win (Section 33.9†).

> 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.

Standing inversion is the failure this stack itself creates (Section 5.2†). The agents that use it accumulate wallets, reputations and collateral, and the human claim above them can quietly stop binding. It fails by key, when no person holds the root, the agent or its operator can raise its delegation ceiling and revocation is only a request. It fails by obedience when a person holds every key and seldom overrides, because a bad override is filed against the person and a bad standing instruction against the instruction. Every board reads green, and the participant the thesis exists to protect is not in the picture. This is a statement about instruments, not likelihood: nothing ranks standing inversion against the other four, and the only claim is that it is the one currently unmeasured.

A design that defends against only the first is defending against the least likely.

5.2 Declining to Energize

Every class so far assumes the machines are running. A state that finds it awkward to ban verification can simply decline to energize it. Curtailment, a queue that never advances, punitive tariffs for compute load, transformers allocated by policy rather than price: none needs hostility toward cryptography, most survive judicial review comfortably, and there is no order to comply with, and therefore nothing to resist. A county board can do the same site by site, and dispersion blunts its veto as it blunts a sovereign's queue (Section 5.5†). The danger is structural. Cheap verification is exogenous to token price but not to the conditions under which power and reference hardware can be obtained. An adversary who can move the physical cost of verification can therefore reach the thesis's primary failure condition without touching the mathematics, the governance or the market. Verification needs obtainable reference hardware at the edge and proving needs bulk power, so bulk electricity prices alone do not break "anyone can verify" (Section 17.10.1†). The response is the stack's usual one: make the exposure measurable (Section 17.8†), price it into Work Credit issuance (Section 25.6.5†), and watch it with the energy-sovereignty red line (Section 17.6). What the stack cannot do is engineer its way out of physics: where third-party proving will not be energized, the answer is dispersion and honest reporting, not a claim that cryptography makes the grid irrelevant.

5.3 Boundary Conditions

Three tensions inside the thesis are stated as limits rather than explained away.

Privacy against coercion. If disclosure becomes mandatory at the chokepoints that matter (exchanges, employers, landlords, visa offices), "optional disclosure" becomes fiction. The protocol has no backdoor or master key, but it cannot stop a holder from being coerced into sharing a key. Non-custodial routes exist, but if every practical route requires identification the option may be theoretical, and if every corridor is choked at once, exit narrows. The protection is technical, and a sufficiently powerful adversary can still coerce the person who holds the key. What it buys is real and bounded: each disclosure needs individual coercion, which makes mass surveillance expensive; optionality survives in less-coerced places; and coercion shows up in telemetry instead of disappearing into platform logs. Telemetry makes coercion visible. It does not prevent it.

Verified compute against the hyperscalers. Decentralized verified compute will not beat the hyperscalers on raw cost; it competes on verifiability, permissionlessness and censorship-resistance. In a benign environment the hyperscalers win on price. Where a provider can be compelled to drop a workload, the value proposition is not "cheaper" but "still available"; where verifiability is the product, the receipt is the value, not the raw cost. The thesis bets that demand for those trust properties is structural and growing, and it is a bet. Whether a prover

sees the input it proves is a separate threat model (Chapter 10†). Beneath the cost question sits a political one, whether the AI frontier is homesteaded or enclosed, which the Homestead Ratio reads on the supply side (Section 5.6†). A frontier can be homesteaded on the supply side and still have one customer.

Holdability against native function. Regulated wrappers may accelerate access, liquidity and legal clarity, and they can also produce holders without users: a fund that holds the asset while shielding its investors from self-custody, native fees and every service the asset exists to buy. The reflex that every ETF is an adversary is as lazy as the reflex that every ETF is adoption. The question is empirical (Section 5.6†): does wrapper growth seed native use, or permanently substitute for it?

None of these defeats the thesis; each marks where it stops. The stack does not promise to solve all social and political problems. It promises technical guarantees that are real within their scope, telemetry that makes a boundary violation visible, and a design that widens that scope over time.

Chapter 6

The Delegation Economy: Who Buys Verification

The buyer of cheap verification is changing shape faster than its supply, and the thesis elsewhere names that buyer only in passing. Chapter 9 and the layer chapters (Chapter 17† through Chapter 25†) specify how Privacy, Proofs, and Compute are supplied, and Chapter 13 how their value might reach a base asset; the argument here is 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 9 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.

6.1 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. What is changing is the marginal buyer of verification, the one whose demand grows fastest and whose behavior sets the price at the margin. It 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 9.1) 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 11. The thesis declines to make the duration warehouse a layer (Section 3.3†),

and the market-realization plane a layer (Section 13.7†), 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.

6.2 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. It also 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 that carry the scarcity: receipts, policy constraints, censorship resistance, priority. 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.

6.3 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.

The line between mesh and fleet is Coase's boundary of the firm [Coase 1937], with verification cost among the transaction costs that set it. Barzel (1982) gives the measurement-cost form: when the attributes of what is traded are costly to measure, markets grow arrangements to economize on measuring them (warranties, packaging, ratings, certification) that would be superfluous if measurement were free. $r(W)$ is that measurement cost for machine work, and the fleet is the arrangement it calls for as the cost approaches that of the work itself.

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 17.6) 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 8) 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.

6.4 The Buyer Bifurcates

Section 7.4.1 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 6.3 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.

6.5 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. 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. A balance authorized in advance lets a delegation pay inside its window without a fresh acquisition round trip (Section 7.1), which is an operating benefit before it is a monetary one. Pledgeability without due diligence (Section 15.1) 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 15.3) 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 14 stays empty, 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.

6.6 Standing

The participation line (Section 3.3) 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, through persistent wallets, reputations built from receipts (Section 20.5†), collateral positions, and 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 11 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 concedes that fee capture at the waist is a discounted-cash-flow quantity (Section 13.2) and that the fork is an equilibrium (Section 13.2.6). 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.

Part II
The Alternatives

Chapter 7

Why Not Something Already Here?

Any case for a new holding has to be made against what a holder can already do. This chapter sets the thesis against the strongest alternatives first (Bitcoin converted on the day, the services bought without the asset, the price owned through a wrapper), then against the objections from what is already here.

7.1 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

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

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

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

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 7.1 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 7.1: Pre-registered stress states for the preholding mechanism. Hypotheses to be tested, not observed cases.

The last row gives the delegation buyer of Chapter 6 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 10), and the contract, not possession of the payment asset, defines the obligation. 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 9.2) 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 18 (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 17.6). 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.

The access experiment's arms are built at equal assurance, so that what separates them is the price of control and exit rather than of verification (Section 34.2.1†). Converting on the day is the holder's alternative. The user's is simpler still: never touch the asset at all.

7.2 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: 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 13.2.7).

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 7.1). 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†).

7.2.1 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 16.3), 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 com-

pute 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 15.2 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).

7.3 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.

The returns on an easily imitated technology may accrue to whoever owns its complementary assets rather than to its developers [Teece 1986], and for an open protocol those assets include the custody, brokerage, and distribution that wrappers already own, the same pressure the fork applies from inside the protocol (Section 13.2.6).

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. It also 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 17.6 and tracked on a dedicated Wrapper Dominance Board (Section 26.9†).

7.4 Bitcoin, Gold, Toll Roads and the Starting State

Those were ways around holding the asset. The objections that follow come from what is already here: Bitcoin does the job, gold needs no network, a toll pays a dividend, and a perfect ledger would harden whatever distribution it starts from. Each is stated at full strength before it is answered, and several are conceded in part.

7.4.1 "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 concedes most of the ground here: Section 13.2.7 states that fees, burns, and collateral establish a cash-flow claim and a balance-sheet lockup rather than moneyness, and Section 15.2 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 15.2 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. Is it sized? No.

- **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†), and 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 16.2). 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 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 15.2 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.

7.4.2 "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.

7.4.3 "Why not just own the toll roads?"

Objection: The rational answer to debasement is not a new money but a claim on a durable toll: a scarce, controllable position (a rail, a chokepoint, a licence, a distribution channel) whose revenue grows with the nominal value of the activity passing through it, without a matching capital investment [Manidis & Roberts 2026]. Such a claim pays a dividend, has a long public record, and repression inflates its revenue along with everything else. A neutral holding pays nothing. The thesis itself argues that durable fee extraction from commoditized throughput makes a utility equity rather than money (Section 13.2.7). An allocator protecting purchasing power should therefore own the utility equity and skip the money.

Response:

- **Across a cycle, the toll may well earn more, and nothing here claims otherwise.** The thesis places durable fee extraction in the utility-equity class and does not argue that a neutral holding out-earns one. The question is what each holding does in the states the hedge is for.
- **A toll is a claim on someone else's throughput, and it carries two exposures a neutral holding does not.** Its rent has claimants: the regulator, the ratepayer, and, at the limit, a sovereign whose backstop entitles it to a share of the gains wherever they accrue, which is the pathway Section 5.2† calls enclosure by rescue. And its value sits inside the collateral loop (Section 3.2†), in which equity valuations are part of the sovereign's fiscal circuit.
- **The two exposures bind in different states.** Under repression or a competent closed-stack stabilization (States 3 and 4 of Section 16.3), the toll can be defended, but on the sovereign's terms: equity or warrants, procurement conditions, and rate regulation decide how much of the rent the holder keeps. In the inverted crisis sequence (Section 3.2.3†, traced in Appendix I.1†), the toll falls with the collateral it is part of. There, the only assets that behave differently are "those whose value does not derive from someone else's balance sheet." A toll's value derives from its payers' balance sheets and its regulator's forbearance.
- **This is an argument for neutral holdings, not for the base asset in particular.** Bitcoin and gold share the property, and the response ranks the base asset above neither: Section 7.4.1 and Section 7.4.2 stand, and so does Red Line 16. Nor is a neutral holding exempt from the political claim. Bans and sanctions (Section 17.3) and the 1933 gold order (Section 29.1†) concede that. The claim is narrower: a neutral holding has no rent for a regulator to reprice.
- **What would refute the response.** If, across repeated episodes of the inverted sequence, neutral holdings move with toll equities and the collateral complex rather than diverging from them, the response is wrong and the objection stands. This is a market reading rather than a red line, for the reason Section 7.4.1 gives for Service-Good Realization: a price-comparative clause cannot be a red line without breaking Corollary 13.1.

7.4.4 "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 3.3.

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 3.3) 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.
- **The channel that does not run through a salary.** Capacity-earned collateral matters more if wages weaken, and early evidence points to weaker entry-level hiring: employment of workers aged 22–25 in AI-exposed occupations stands below where it would otherwise be [Brynjolfsson et al. 2026]. A newcomer who supplies verified work earns the base asset directly, without first converting a wage. The buyer map's operator series reports that channel with a new-participant split (Section 26.8.8†), classified by principal (Section 26.8.9†) so that agent fleets do not count as entrants. It is reported beside the refuter below and does not replace it: newcomers who can earn the asset but not buy it from income still trip that refuter if incumbent concentration rises. Adjusting it for a falling labor share would loosen

a falsifier to fit the facts.

- **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.

Chapter 8

The Closed Sovereign Stack

Everything argued elsewhere in this thesis is argued in the abstract: when soft guarantees weaken, value migrates toward hardened, verifiable, vertically coherent capacity, and the layer chapters (Chapter 17† through Chapter 25†) describe how to build such a stack.

Is anyone actually doing this, and does it work?

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 17.6 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.

The case is China's: an energy and industrial strategy best read as a sustained bet against the reliability of external guarantees, which builds many deliberately redundant paths from a controlled input to a required output and accepts poor measured returns in exchange for freedom of action (Section 33.1†).

8.1 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.5† 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: a profound regime change, though not 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 3.3 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.

Read "verify" as "retain an independent means of knowing", and the closed stack's instruction is the cypherpunk one: *don't trust; possess, duplicate, stockpile, or retain an alternative conversion pathway* (Section 33.2†).

8.2 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 this comparison is usually drawn. The convenient framing sets a closed foreign stack against an open Western one. The evidence in Section 8.1 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: whether a participant may exit without permission.

Recourse, however valuable, is not exit.

The case also makes claims in this thesis harder to dismiss as utopian: hardened stacks are

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 6, Corollary 30.1†
Chief danger	Enclosure by rescue	Administrative enclosure	Fragmentation and weak value capture	Section 5.2†, Lemma 13.1

Table 8.1: Three sovereign stacks: convergent structure, different locus of sovereignty. The first two columns differ in method and agree that the state holds the keys.

achievable, serious actors price resilience, and the physical layer is monetary (Section 33.6†).

8.3 What This Warns

The warning is sharper than the validation, and it is the reason this chapter sits in the argument 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 3.3 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. That is a claim about key custody, exit rights, disclosure defaults, and value capture, not about throughput.

Part III
The Instruments

With the alternatives faced, the thesis has to say exactly what it proposes, and on what conditions. This Part does: the triad and the chain it would have to complete, the asset and the rules that route value to it, the designs under test, and the Work Credits kept apart from the asset.

Chapter 9

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** (a user can check whether they 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.

9.1 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 the whole thesis, 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 6: 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 6.2); 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 17.6).

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 13.3, Section 7.2).

Severed if: users obtain equivalent triad capacity through hyperscalers, stablecoins, or custodial APIs without the asset—Red Line 6 (Section 17.6, Section 28.2†).

8. **A monetary-risk warehouse forms.** Gross native demand is not a monetary 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 6.5).
   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 14, 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 15.1), denomination (Section 15.3), and the state-contingent holder-side service flow (Section 15.2), each with its own instrument and red line. The third must first beat its strongest alternative, Bitcoin converted on the day, and the access experiment tests whether it does (Section 7.1).
   Severed if: any of the nine store-of-value requirements fails, telemetry showing them is captured, or any of the three mechanisms reads null on its instrument, as its red line specifies; if only outsiders decline to reckon in the unit, the link is weakened rather than severed.

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 the thesis argues: acceptance without due diligence (Section 15.1), reckoning in the unit (Section 15.3), and the state-contingency of differential value (Section 15.2). Does the thesis size them? No. Two questions stay open beside them: whether the convenience yield, if it exists, exceeds the premium gold or Bitcoin already carries, and whether the open duration warehouse (Section 34.5.2†) can finance long-lived infrastructure without concentrated credit intermediaries. Ratings that go up are self-congratulation; ratings that identify the unmeasured bridges carry information.

Four questions, and where the answers branch. Read from top to bottom, the nine links look like a procession in which each success hands the next its premise. The thesis already allows every branch: the stack is fundable as public infrastructure whether or not a coin follows (Section 18.1), Work Credits remain service claims however well they perform (Section 12.4†), and the architecture comparison leaves open whether the three services share one asset at all (Table 8.3†). The links answer four separate questions (Table 9.2), and a yes to one does not carry a yes to the next. Services can succeed without the asset capturing their value; the asset can capture value without offering its holders anything a buyer on the day lacks; and a held asset can attract demand that never becomes the self-custodied, agency-preserving use the ninth requirement asks for (Item 9). Each exit in Fig. 9.1 ends in an outcome the thesis already names.

The procession is one path through a tree.

What follows defends links 7 through 9, where the monetary argument actually lives. The stack chapters defend links 2 through 6 by building and measuring what supplies them

Link	Strength	Why
Utility demand	Strong	Structural drivers are already observable; the delegation mechanism (Chapter 6) 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 13.2.7)
Non-bypassability	**Medium**	Written as a binary; the economics requires a magnitude (Δ) that is nowhere estimated (Section 13.2.4)
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 15.1, Section 15.3, Section 15.2)

Table 9.1: 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.

Question	Links	Outcome if the yeses stop here
1. Useful services. Do buyers keep paying for private settlement, portable proof, and verified compute through the cycle?	1–5	Useful infrastructure without money: the services are fundable on their own terms, and their fees may settle in fiat or stablecoins rather than the asset (Section 18.1, Section 16.1.2)
2. Asset value capture. Do the fees accrue to a native asset that demand cannot route around?	6–7	A claim on a fee stream and a collateral lockup, both quantities a discounted-cash-flow valuation captures (Section 13.2.7); the service-asset reclassification of Warning 13 (Section 30.7†)
3. A distinct holder service. Does holding the asset deliver what buying it on the day cannot?	9 (holder-side flow)	A balance worth holding inside the band where conversion fails and native settlement survives (Section 7.1, Section 15.2); pledgeability and denomination are read separately, each on its own instrument
4. Monetary use consistent with the mission. Is the asset pledged, reckoned in, and held by a self-custodied, loss-bearing constituency, in a way that preserves agency?	8; 9 (pledgeability, denomination)	The end of the chain: a defensible monetary candidate, from which nothing follows about price (Corollary 13.1). Demand met through wrappers and custody is not a yes here (Section 7.3)

Table 9.2: The conditional chain as four questions. A yes to one question does not carry a yes to the next, and each row names where the chain ends if the answers stop there. Link numbers follow the enumeration above; the strength of each link is rated in Table 9.1.

(Chapter 11), and the telemetry and the red lines make every link falsifiable (Chapter 26†, Section 17.6).

9.2 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. What determines the service is the topology: the transformation path, location, actor, and horizon.

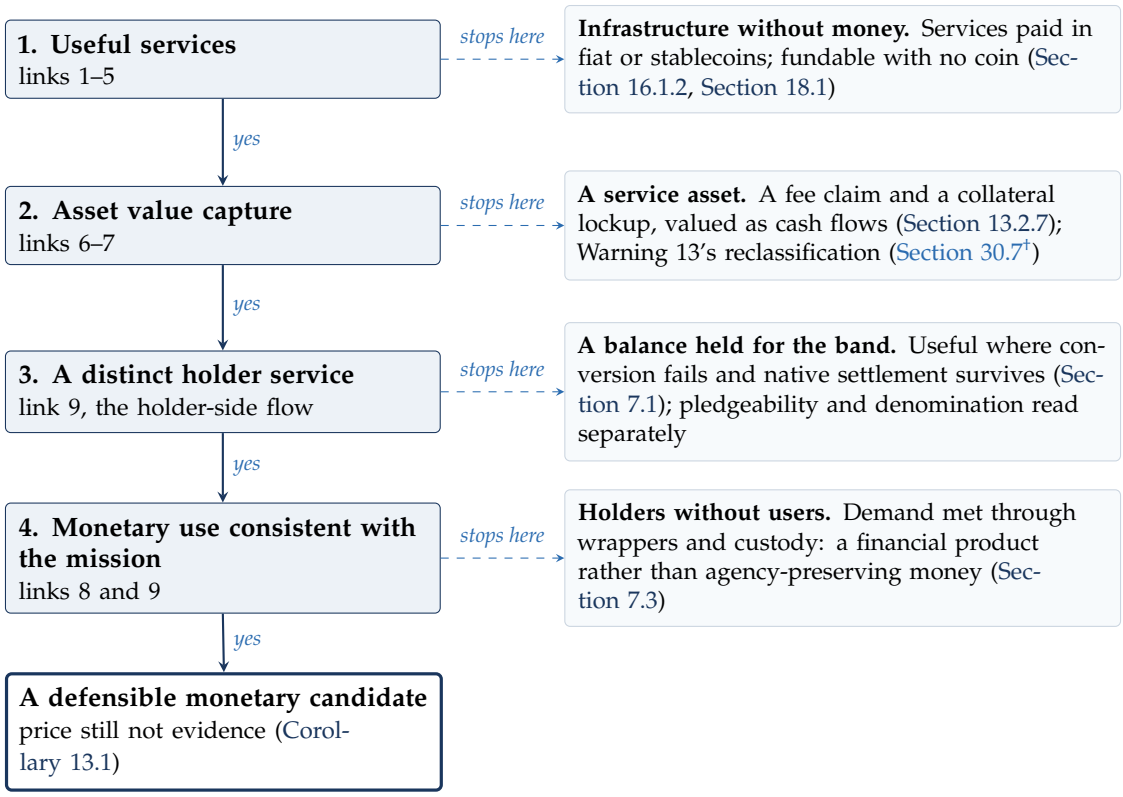


Figure 9.1: The chain drawn as branches. Each question can be answered yes while the next is answered no, and every exit to the right ends in an outcome the thesis already describes. The path down the left edge, a yes at every question, is the nine links completed.

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.

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 6.4 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. $A_{\text{full-stack}}(R)$ is read the same way, because redundant paths under one controller, creditor, or jurisdiction fail under the same pressure and count as one path (Section 26.8.10†).

9.3 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?**
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 9.3: 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, trans-

ferable, 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.

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 there are 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 13.2, 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 13.2.7 shows why that reading fails, and Section 15.2 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). Even where they cannot, the demand it creates is for a *medium of settlement* held for seconds, not for a store of value

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 9.4: 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†).

held through a cycle. Section 13.2 bounds what that toll can be, and Section 7.2.1 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 13.2.7:

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 9.4), 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 13.2.7 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 taken up in Section 13.2 and answers in the negative at Section 13.2.7.

9.4 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 13.2.7). 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 7.3). They matter because Section 13.2 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 10

Work Credits: Reserved Service Claims

Whether there is a fee to route belongs to the case for holding. First, the instrument the designs keep apart from the monetary candidate: Work Credits, typed claims on service. This chapter defines them, separates them from the receipts that prove work was done, and asks what would have to be true for one to become money.

10.1 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: one proves the past, the other reserves the future. 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.

10.1.1 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 a 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.

Why a Work Credit is not savings is physical, not rhetorical: exposed to basis, location, hardware, expiry and delivery, it is a warehouse receipt with a date on it, and a savings instrument has to survive the day the warehouse is the problem (Section 12.4†).

10.2 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, and balances retained through drawdowns instead of recycled into service. The tell that would settle it is 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 15.2).

Part IV
Delivery

Every instrument in the previous Part assumes a stack that keeps delivering when it is attacked. The threat model named the attacker, a rational sovereign under stress reaching for the levers it controls (Chapter 5).

Chapter 11

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.

The architecture is therefore worth making 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 layer chapters (Chapter 17† through Chapter 25†) 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. The layer chapters (Chapter 17† through Chapter 23†) show that the stack *can* supply the triad. Chapter 9 and Chapter 13 must show separately that a specific asset *captures the value* of that supply. Without the value-capture bridge developed in Chapter 9, 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 (Section 18.3†) 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 chapter (Chapter 19†) 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 (Chapter 20†, Chapter 27†) 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 6.6). 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 thesis's economic core. 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. Verification Asymmetry & VerifyPrice (Section 22.2†), Proof Factories (Section 22.5†), AI-PoUW (Section 22.4†), and the modular stack (Chapter 24†) describe 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 (Section 23.2†, Section 23.3†), and the communications and identity chapters (Chapter 18†, Chapter 20†), 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 layer chapters (Chapter 17† through Chapter 25†) 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.

11.1 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 condi-

tioned 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.

The full design turns this into a "substrate kit" of twelve primitives and four reference applications (Section 24.1†). Outside Section 24.1†, the names PAL, PIDL, privacy-rails kit, proof factories, and VerifyPrice observatory serve as shorthand for these building blocks.

Chapter 12

How the Stack Delivers

The layer outline says where each piece of the stack lives. This chapter asks two plainer questions: what does each layer hand to the one above it, and how would anyone outside know that it did? The engineering that answers the first lives in the full edition. What a reader needs here is the deliverable, the number that measures it, and the reading of that number that would show it failing.

One limit governs everything that follows. The layers create markets for proofs, verified compute and private settlement; they do not create money at all. A working market for verified FLOPs is necessary to the monetary thesis and not sufficient for it, and the Work Credits that ride on these layers remain typed service claims rather than monetary instruments. Whether any of the stack's economics accrue to the base asset is the value-capture question (Value Capture Lemma); whether accrual becomes moneyness is the separate question of the three link-9 mechanisms. Delivery is the floor those questions stand on.

12.1 What Each Layer Hands Up

The lower four layers answer four questions. Can the machine be trusted? Can it be reached under censorship? Can code get onto it and stay current? Can an actor prove who it is without being doxxed? Their deliverables are matter that can be audited, packets that survive pressure, code that ships without a gatekeeper, and identity that travels without a name. Layer 0 does not hand up a promise about hardware. It hands up artifacts: hardware profiles with their sampling history, attestation receipts binding a device to a profile and a firmware hash, streams of power and health telemetry, and incident flags of the kind "this profile is compromised; accept no new work" (Section 17.13†). Each of the four was specified so that its failure is a measurement, not a debate (Section 20.12†).

The economic pair comes next. Layer 4 turns computation and provenance into standardized, verifiable work units: proofs and verified FLOPs that any counterparty can accept without trusting a platform, because checking them is cheap and public. Layer 5 turns those units into value flows that are non-custodial, refund-safe and private by default, with disclosure at the holder's discretion (Section 23.2†). Every party to a failed swap has a no-loss exit it can take alone, within bounded time (Section 23.4†).

Layer 6 holds the rest to account. Most protocols ship documents called constitutions. In practice the real constitution is whatever cannot be silently violated without getting caught, and for this stack that means a set of service-level objectives on capacity and neutrality, plus the dashboards and receipts that make a violation obvious (Section 25.1†). Those objectives are not operational goals with a monetary flavour. They are the monetary invariants.

Built, and not yet built. A reader asked to believe the stack can be built should be told which parts already can be. On the proving side, generic succinct proofs, Freivalds-audited matrix products and inference checked by sampled re-execution are feasible in production; succinct proving of large matrix products is research, and so is zero-knowledge proof of large-model inference (Section 22.6†). The consensus security of MatMul-PoUW is a conjecture, and inference PoUW is pilot-stage. The research items are published as registry targets so that the destination is public, and a maturity gate keeps every workload from issuing full-weight Work Credits until its verification component reaches production maturity (Section 22.10†). The gate lifts only when the maturity table lifts, and the maturity table is moved by deployment evidence, not by argument.

12.2 Verification Asymmetry

The spine of Layer 4 is a ratio. For a workload W, take the cost to produce a result and its proof, $p(W)$, and the cost to verify them, $v(W)$, both priced in the same numeraire; then $r(W) = \frac{v(W)}{p(W)}$ is the **verification asymmetry**, and the layer's whole job fits in one line:

> **Make $r(W) \ll 1$ for the workloads that matter, and keep it that way in production.**

This is the gold touchstone carrying a price: expensive to smelt, cheap to assay, and $r(W)$ is the assay, priced. When checking is small and stable, anyone (small nodes included) can check, and proofs and verified FLOPs become commodities rather than favours. The ratio has a trap built in, so it is published with both components. It improves when verification gets cheaper and also when production gets dearer, and a passing ratio bought with expensive production is a cost problem wearing an asymmetry credential.

A single ceiling on $r(W)$ does not protect the property the monetary claim needs. At a ceiling where checking costs a sizeable fraction of producing, nobody checks everything: participants sample, delegate, or take a counterparty's word, which is the symmetric ignorance of Section 15.1 reached through an SLO the network is passing. So verification is sorted by *modality* (Section 22.2.1†). A succinct proof (M1) is checked at a cost that grows more slowly than the work proved, and usually not at all. An algebraic or probabilistic check (M2) is far cheaper than re-running the work and states its soundness error alongside the ratio. Replica-

tion or attestation (M3) buys fault tolerance and accountability rather than asymmetry, since its cost stays on the order of doing the work again.

Only M1 and M2 carry the monetary mechanism, and this is a consequence rather than a preference. Under M3 the participant is not checking; a committee is, and a lender asked to take a position on a committee's word has to form a view of the committee, which is due diligence, the very thing information-insensitivity is defined as the absence of. M3 is the right construction for a great deal of useful capacity. It supports service delivery, Work Credits and SLA tiers. It does not support pledgeability.

> Replication buys fault tolerance. Succinctness buys moneyness. The thesis needs both and has been pricing them as though they were the same good.

Modality is also what a collateral haircut was always pricing. An M3-verified position carries delegated-trust exposure that no hardware grade removes, and an M1-verified position on modest hardware may carry less, so the haircut schedule is indexed on modality and grade jointly rather than on either alone (Section 17.4.2†).

Three reporting rules follow, and none relaxes anything. VERIFYPRICE series are published banded by modality, because an aggregate across a mixed portfolio can improve by changing the mix rather than improving anything. The M1+M2 share of verified capacity is a published series of its own, beside fee coverage on the Value Capture Board, since it is the fraction of output that carries the monetary mechanism at all. And Red Line 3's monoculture condition gains a second reading: a capacity mix migrating from M1 toward M3 is a quieter form of single-vendor concentration, since both end at "trust the verifier," so that share carries a published floor and a fall below it is Sev-1. The rhetoric can then be correct: any claim that checking costs a tiny fraction of producing must name its modality wherever it appears, because it is true of succinct verification and false of replication.

12.3 VerifyPrice

Verification asymmetry is a design goal; VERIFYPRICE turns it into a dashboard. It has exactly one definition, a five-field tuple per canonical workload: median and ninety-fifth-percentile verification time, the same two percentiles of cost, and the failure rate, measured on reference hardware (Section 22.6†). Everything that resembles it elsewhere is an input to it, not a rival.

The harness that produces the tuple starts from an uncomfortable assumption. The observatory is not a disinterested instrument; it is a publisher with interests, run by operators who can be pressured, and its numbers are exactly the kind a protocol under strain would prefer not to publish. So the specification assumes the observatory is adversarially fallible and is written so that when it 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.**

The targets come in three tiers, and their order matters (Section 22.7†). The constitutional tier is measured in real resources on reference hardware and is exogenous to token price: whether the token rallies or dumps, it must hold, and a breach is a Sev-1 incident that requires remediation. The operational tier asks whether verification stays affordable relative to typical fees, and it carries the same trap as the ratio, so an improvement achieved on a rising fee is reported as neutral, not as a pass. The market tier quotes verification in tokens. It is a useful signal and never a target, because token price is endogenous and a target built on it would be circular.

Exogenous to token price is not exogenous to everything. Real-resource targets depend on whether unprivileged reference hardware stays obtainable and on what power costs at the edge, and a state can move both. Red Line 1, the thesis's hinge, can therefore be tripped from outside, which is why it is watched upstream, through sovereign optionality and the energy-sovereignty red line, rather than left to be discovered at the hinge (Chapter 5). There is a deeper reason to insist on the target, if the framing of Section 2.0.1 is right that a medium restructures only those who habituate to it. Verification must then be cheap enough to become a habit before it can become anything else, and a cost that tracked the token's price would keep the habit from forming.

A proof can be cheap to check while the facility behind it is not. **Physical VerifyPrice** prices that second check: what it takes to confirm what a facility says about its power, its hardware and its jurisdiction (Section 22.8†). A Work Credit cannot be high-assurance if verifying the proof takes seconds and verifying the facility takes months of trust-the-vendor diligence, so a facility whose Physical VerifyPrice rises above threshold sees its credits issued at reduced weight or capped.

12.4 Delivered Verified Capacity

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, and between the two runs a network of mains, valves and pumping stations, any one of which can decide whether water arrives. Delivered Verified Capacity is the tap, not the reservoir.

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.

Because it is computed as a flow rather than a sum of capacities, the binding constraint, the narrowest main, is the only number that matters while it binds: the sum flatters, the cut confines (Section 17.9†). Three features keep it honest. It is specific to the scenario, because a drought, a fuel shock, a transformer loss, a chip embargo and a network partition are different graphs, and an average across them hides the edge that fails. It counts substitution latency, because a backup that takes three weeks to route is not redundancy for a fire that started today. And it assigns dependency groups, because facilities that share a transformer vendor, a firmware signer, a cable or a stablecoin can fail together, and geographic dispersion without dependency-group dispersion is not independent capacity.

DVC is where delivery binds issuance. Work Credits and any other forward service claim are bounded by the worst pre-declared scenario's DVC over the issuance interval, under a published haircut. Facility and proof receipts supply the evidence for each edge's capacity; they are not capacity themselves. Governance may not replace this bound with nameplate power, installed hardware, gross proof throughput or a benign-state average. A scenario in which nothing is delivered binds issuance at zero, unless the constitution has declared it force majeure in advance, and that is the correct answer.

Two limits belong beside it. DVC is reported next to sovereign optionality rather than folded into it: a network can have many small, slow alternatives and little delivered capacity, or ample capacity on one large path and few alternatives, and a reassuring composite would hide either. And a rate is not a joint completion. Delegated work arrives as trees of units that must all complete inside one refund window, so a network can deliver its DVC in units per interval and still fail most deep trees during a disruption. That gap is measured (Section 22.9†), and in this version no issuance bound is tightened on it.

12.5 Telemetry Anyone Can Re-Run

Every number above is only as good as the regime that produces it. The most common attack on telemetry is not fabrication but curation: nobody has to invent a number when selective reporting will do, and a dashboard fed from good regions, fast hardware and favourable conditions flatters itself without a single lie in it. The telemetry families exist to make that detectable.

Reference hardware prices "anyone can verify" on hardware a person can actually own: constitutional targets are defined on a laptop-class machine unless a workload's registry entry names another class (Section 26.1†). **Independent implementations** turn a silent bug into a visible disagreement, because a single verifier codebase is a single point of failure for the truth

itself, and a disagreement, unlike a bug, shows up on a dashboard (Section 26.2†). **Sampling** is seeded by public randomness and stratified, and a dashboard point without a verifiable seed is not part of the record (Section 26.4†). **Data availability** keeps the bad months from vanishing: receipt sets are content-addressed and their roots committed to a neutral ledger, which does not prevent deletion but makes it visible, and visible deletion is itself a datum (Section 26.5†). **Reproducibility** closes the loop. A datapoint is a dataset root, a harness version, a verifier version and a result, and anyone can fetch the receipts, run the same versions and compare; a mismatch is a Sev-1 incident (Section 26.6†).

Underneath them sits one discipline: **if it matters, there is a receipt for it.** A claim about performance, neutrality or safety that cannot be tied back to receipts anyone can re-check is out of scope for the monetary thesis (Section 26.7†). Three primary dials read the stack itself. Is verification still cheap? That is VERIFYPRICE. Can users still reach the network under pressure? That is VerifyReach. Can they still move value privately and safely? That is VERIFYSETTLE. A fourth family, VerifyFlow, reads the market outside it: how the asset is represented and priced there.

These feed public boards (Section 26.9†). The Value Capture Board asks whether service value is accruing to the native asset. The Market Realization & Wrapper Board, which reports VerifyFlow, asks what is currently setting its price. They must never be merged: collapse them and the distinction the thesis depends on disappears, because a price can be set by forces that have nothing to do with service value reaching the asset.

Telemetry also turns neutrality and repression-resilience, usually sold as moral claims, into hypotheses (Section 26.11†). Neutrality says no actor or group can reliably censor, privilege or front-run specific flows across Layers 1–5 without being detected and routed around. Repression-resilience says that under yield-curve control, capital controls, deep packet inspection and blacklists, the triad keeps enough delivered capacity, and the base asset's self-custodied holders enough agency, to remain a monetary candidate. If flows stay blocked with no alternative emerging, neutrality is falsified; if privacy corridors collapse under regulation, repression-resilience is; and if verification cost is pushed beyond reach, so is "anyone can verify." **A system that cannot tolerate bad news cannot be money; it can only be marketing.**

The stress harness publishes, for each policy attack, the test and the pass condition in advance (Section 26.12†). The market around the asset gets scenarios of its own, because the compositional adversary responds to none of those tests. One of these scenarios looks like good news: price triples and wrapper holdings grow while native fees, settlement, proofs and compute stay flat. The expected result is market success, monetary-adoption failure and an automatic warning on the dashboard.

A thesis that cannot raise an alarm during a rally is not falsifiable.

Part V
Why Hold

Chapter 13

Why Hold the Base Asset?

A holder who wants verified service in a bad state can hold a Work Credit, hold Bitcoin and buy on the day, or hold the base asset. The case for the third was first set against Bitcoin (Section 7.1). This chapter follows the channels through which service value reaches the asset, and asks how much of that value is money.

13.1 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 13.2.7) 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 9.3). 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 13.2.4), 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.

The link needs two precisions. 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 13.2.3). 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 15.1). 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 15.1 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 15.3). 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 15.2, 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. Is either sized? Not 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 13.2 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 15.2, 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 13.2.7, Section 15.2).

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 15.1. 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 13.2 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 7.1 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 13.2) or while accrual occurs and no monetary premium follows (Section 13.2.7). □

Which conditions the monetary channels actually use. Of the five conditions, four govern *accrual*, which is a discounted-cash-flow quantity. **Issuance discipline is the exception**: a discretionary supply is information-sensitive on supply, for the reason Section 15.1 gives. Condition 4 is therefore doing monetary work that Conditions 1, 2, 3 and 5 are not. That is the

same conclusion Section 15.2 reaches from the other direction, and the two 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 13.2 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 the claim that fees, burns, and collateral earn the asset a monetary premium (Section 13.2.7). 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 9.4. 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 corollary follows from the existence of the wrapper rows, not from any assumption about their size; Appendix H† quantifies the size. □

13.1.1 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 15.2.

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 13.1: 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.

13.2 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 res-

cues considerably less than it first appears to, and what survives is not the claim that fees, burns, and collateral earn the asset a monetary premium.

13.2.1 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 13.2.3; 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 13.2.3 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 reprices 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.

13.2.2 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 13.2.3.

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 17.6) 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.

13.2.3 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 a red line measures it (Section 17.6).

13.2.4 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 9.1). 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.

13.2.5 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 13.2.3 covers four properties: cryptographic checkability, neutrality, credible non-discretion, and non-custodial exit. Neu-

trality 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 13.2.3 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.

The tail is one limit on the fee; openness is another. Every property the thesis is proud of, from published workload specifications to permissionless participation, also lowers the cost

of copying the protocol (Section 13.3.7†).

13.2.6 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 13.2.3 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 13.2.3, which stated only the first term, and it is the correct place to put the concession Section 13.3.7† made informally.

Contestable markets and coordination games give opposite predictions, and only one of them is the right model. The objection that opened Section 13.2 (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 but 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. Section 13.3.7† 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. 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 17.6 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 15.3 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 9.1 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.

Openness has one further effect that neither bound prices: a protocol anyone can implement is a credible entrant against the incumbents it contests, and the credible possibility of entry can discipline conduct without any entry being observed [Baumol 1982], which can serve agency while leaving the asset nothing to capture (Table 9.2).

13.2.7 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 13.2.3) 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 13.1 are design commitments; nothing in this thesis claims that they set a price. The incidence question asked of the fee is asked of each monetary mechanism's benefit as well (Section 15.1).

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 7.4.1.

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 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 restated 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 15.2. The Value Capture Lemma formalizes the accrual conditions; the telemetry regime (Chapter 25†) makes them falsifiable.

Part VI
When It Would Be Money

Fee accrual buys a cash-flow claim and a collateral lockup, and on its own nothing more. This Part asks what else would have to be true for the asset to become money, and how anyone would see it happen.

Chapter 14

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 6.5). 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, and the holder-side stress-deliverable service is the medium. That 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 9.1 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.

14.1 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: 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.

14.2 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.

14.3 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. VerifyReach 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.

14.4 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 7.4.1). A store-of-value premium, where it exists, accrues to assets that preserve agency on whichever surface is under attack.

Chapter 15

The Three Monetary Mechanisms

The case for the possibility is made; what remains is where a premium would actually come from. The answer is three mechanisms on the chain's last link: pledgeability without due diligence, the state-contingent holder-side flow, and denomination. Each is argued here and read on its own red line (Section 17.7).

15.1 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.

The distance between those two is what the study of collateral and short-term funding means by moneyness. Not scarcity. Not cash flow. Not convenience yield. 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. A funding run is not a change of mind about fundamentals. It is the moment looking starts to pay.

Two kinds of information. Transparency alone produces nothing of the kind, since equities are the most transparent securities in existence and among the most information-sensitive. The claim holds for one kind of information only. *Fundamental* information is what the asset will be worth: cheap verification does nothing to it, the thesis claims nothing about it, and a triad asset stays as sensitive to it as gold does. *Unit-quality* information is whether *this* unit, *this* receipt, *this* pledged position is what it purports to be.

Why did a market full of experts stop accepting collateral at face value in 2008? Not because anyone changed their view of housing; every participant had held one for years. What nobody could answer cheaply was whether a *particular* pool was good, so lenders began producing private information about specific collateral, haircuts widened, and no-questions-asked acceptance ended [Gorton & Metrick 2012]. Moneyness failed on unit quality while fundamentals were already common knowledge. The same split says where collateral demand

would lean if private views on fundamentals grew more valuable: toward public unit quality, the one axis on which nobody can hold an edge.

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.**

Deterrence, not audit. The box does not claim that anyone looks. Universal audit is observed nowhere, even where looking is close to free: public filings sit unread, and the record is routinely complete and unanalyzed [Green 2026f]. An argument that needed every counterparty to verify would be refuted on its first day, and would deserve to be.

This argument needs nobody to look: 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 agent who looks acquires nothing the agent who declines could not have had at the same cost. A security camera does most of its work on the days nobody watches the tape. The discipline falls on the claimant, before the claim is made, the way the credible possibility of entry disciplines a contestable market without entry being observed [Baumol 1982]: a misrepresentation anyone 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. What to measure is the *cost* of verification, not its observed rate: Red Line 1 reads the standing threat, and falling verification activity is not evidence against the mechanism while the cost holds. And a system can be perfectly auditable and monetarily inert, because deterrence works only if being caught would cost the claimant something, a

condition that Red Line 14 and Section 13.1 instrument rather than assume.

The position, not the token. That a unit of a public ledger is authentic is true of every cryptocurrency, and nobody 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*. Pledge the asset against an obligation to deliver a service (a prover bond, corridor collateral, SLA escrow) and a lender has to assess the bundle: whether the capacity behind the obligation exists, whether the operator's slashing exposure is what it claims, whether the physical substrate delivers under the scenario the haircut assumes. That is where due diligence lives, and it is what Facility Capacity Receipts, energy receipts, Delivered Verified Capacity (Section 17.9†) and the grade ladder make cheap, which is why Layer 0 is a *monetary* requirement rather than an engineering preference (Chapter 4).

The thesis already prices this. Gold, expensive to smelt and cheap to assay, is information-insensitive on unit quality, which is why a stranger's nugget could be accepted without investigating the stranger; every haircut keyed to Layer 0 grade, sovereign optionality, or bridge class (Section 17.4†, Section 25.6.5†) 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.**

Necessary, not sufficient. 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 grades a stone without making stones interchangeable), **durability and portability** (grain is gradeable, perishable, and bulky), and **the absence of a counterparty** (a warehouse receipt brings back the custodian, 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 new term. Assaying more metal costs more; verifying a succinct proof costs the same whether the claim beneath it is worth a dollar or a billion. If information-insensitivity is what makes collateral money-like, an assay invariant to notional is a change in kind rather than in degree.

The forger gets better tools too. That is the strongest objection to any cheap-verification argument: if fabricating plausible false claims gets cheaper as fast as checking them, nothing has changed. In natural language the objection is fatal, because telling true from false degrades as the adversary invests. Cryptographic verification has no such gradient. 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 but infeasible, and rejecting one costs what accepting a valid one costs. A thousand forged proofs are a thousand rejections, 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 is as firm as the claim (Chapter 10†): a proof is a bounded attestation, not truth, and says nothing about whether the workload was worth running or whether the sentence a human attached to the artifact is honest. The invariance holds only over claims that can be stated as a computation, a class much narrower than the one on which trust is failing, and outside it the thesis has nothing to offer.

Free riding, and discretion. Grossman and Stiglitz [Grossman & Stiglitz 1980] showed that information nobody is paid to produce does not get produced. Proof systems escape the trap because cost and benefit fall on different parties: the prover pays to be believed, and every counterparty consumes the result at negligible cost. Verification is a private good bought by the claimant, and its public checkability cannot be withheld without destroying what was bought. Read from the financing side, competition prices proving toward its cost (Section 22.2†), which is why the fee stream is ordinary service revenue (Section 13.2.7). **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.** Whatever premium exists is not in the proving business, which is why Section 15.2 and Section 15.3 look for it somewhere a valuation cannot follow.

Issuance discipline, Condition 4 of the Value Capture Lemma, belongs here, and not as scarcity. A supply set by discretion is information-sensitive on supply: every counterparty must form and maintain a private view of what the issuing body will do, which is due diligence by another name. **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.**

How it fails. Symmetric knowledge reverts to ignorance when verification stops being affordable, concentrates until most participants trust a verifier instead of verifying, or loses its receipt data. Each reads as a degradation of service, and each is also a **monetary run trigger**. A fourth mode has no red line: **scope drift**. A position legible on its proof and opaque on

its operator's off-ledger leverage, counterparty concentration, or rehypothecation is 2008 with better cryptography. It is published as the off-ledger exposure disclosure series (Section 26.9†), which feeds the Red Line 17 panel and the Red Line 10 opacity reading. The gap is a number, not a caveat.

What it predicts. A reader of Dang, Gorton and Holmström [Dang et al. 2020] will object that insensitivity belongs to the *payoff*: drive information cost to zero on a state-dependent payoff and it becomes *more* sensitive. The objection is correct, and it fixes the scope. The argument here is adverse selection on unit quality, not a claim about payoffs. Bitcoin is the control: its units have been freely verifiable since 2009, and its haircuts sit at 30–50% and vary widely across lenders, because they price realized volatility and liquidation depth, which no verifier compresses. The thesis does not predict lower haircut *levels*.

It predicts something narrower. Quoted by the same lenders at the same moment, against opaque collateral of similar volatility, a position whose backing is publicly verifiable should show lower cross-lender *dispersion*, and dispersion that does not widen relative to the comparables when the market does. The flatter the payoff, the more of the prediction applies: an over-collateralized, DVC-backed position, debt-like over the states the haircut assumes, carries all of it, and the bare base asset a weaker share, since a lender who wants a flat payoff builds one from the asset and a haircut. A funding run's signature is dispersion widening *before* level, as the better-informed lenders leave first; if that signature appears here, against the comparables and with receipts available, symmetric knowledge did not hold.

Is it measured? Every quarter, against a pre-committed panel of lenders: haircut level, cross-lender dispersion as the primary series, the cost of onboarding the asset at a new venue, and dispersion under stress. Is it sized? No. **Red Line 17** reads the prediction as its box specifies (Section 17.7), on a design that lets Bitcoin's native units, set against their wrapped representations, pass or fail rather than fail by construction. Nothing in it says who keeps the gain: the saving goes first to the borrower who posts the position, and whether it travels on to the asset's holders is read from the buyer map's sell-through and retained balances (Section 26.8.8†), a reading with no trigger of its own.

The section claims no price stability, no insensitivity on fundamentals, and no automatic acceptance by any lender, and it raises no rating in Table 9.1. It 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.

Acceptance and demand. Information-insensitivity is why a counterparty would *accept* the asset without investigating it; the holder-side flow of Section 15.2 is why a holder would *want* it when other arrangements fail. Acceptance without demand is a payment rail; demand without acceptance is a collectible. Section 15.3 asks the third question, whether anyone will *reckon* in

it. The thesis argues all three, sizes none, and publishes an instrument for each.

15.2 Where a Premium Would Come From

If fees, burns, and collateral yield a well-valued cash-flow claim rather than money, the monetary claim has to be relocated, not abandoned. It relocates to something the thesis has been describing all along without calling it the monetary argument.

Δ depends on the regime. Write $\Delta(R)$, with R the regime-pressure index of Stage A (Section 16.1.1). Δ is thin in benign states because a solvent, suable indemnitor substitutes for a proof (Section 13.2.5). Turn that sentence over and it says when Δ is large: when the counterparty cannot be sued, will not be solvent when it matters, is itself the adversary, or inflicts a loss no damages restore (censorship, seizure, deplatforming, disclosure). That list is the threat model of Chapter 5, restated as a demand curve.

> 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.

The tempting next step is to say that a payoff arriving when wealth is scarce must be worth more than its discounted expectation. It fails the test already set: a monetary premium is value *in excess of* discounted cash flows (Section 13.2.7), and a valuation run with the correct stochastic discount factor already weights every cash flow by its state price. So the argument splits in two.

A better utility. Because $\Delta(R)$ rises when the substitutes fail, the fee stream is countercyclical, and an investor should pay more for it per unit of expected cash flow. That is a discount-rate effect, fully representable inside a discounted-cash-flow valuation. A countercyclical utility is a better utility, not a monetary asset. Conceded, and set aside: nothing below rests on it.

A service to the bearer. What the holder obtains is not a cash flow, so there is nothing to discount. It accrues to whoever holds the unit, by holding it, and a valuation of distributions omits it however its discount rate is set (one that prices services from possession, as transactional-demand models of tokens do, need not). It is why a bearer instrument held outside the banking system, with no cash flows at all, can carry a price. Storage theory has the mechanism as **convenience yield** ([Kaldor 1939]; [Working 1949]), the holder of inventory owning an option against stockout that the holder of a delivery claim lacks; the analogy is to

the mechanism only, since a triad asset is not a storage commodity and takes no magnitude from that literature.

A fire extinguisher and a contract for the delivery of one both have a price. Only one works during the fire, and every argument here about regime states is an argument about when the market is on fire. An extinguisher is not money, though. It establishes a benefit of ownership; the step to monetary use is independent acceptance and a persistent demand to hold, which this section does not supply (Section 15.1 argues the acceptance).

Gold makes the same point from the other side: no cash flows, negative carry, a storage-theory yield near zero in normal regimes, and a price every constructible DCF falls well below. On the standard account the residual is monetary and portfolio demand, the holder-side category reached from another 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 13.2.7).

What it rests on. A premium carried by the holder-side flow rests on whether a holder can use the thing under pressure, and none of that is fee routing. Of the five lemma conditions, only issuance discipline touches the channel, because credible non-discretion is part of the payoff rather than of the accrual. Link 8 asks for a persistent self-custodied constituency able to bear loss. The premium itself hangs off link 9: liquidity, neutrality, verifiability, legal holdability, and agency preservation (Item 9).

A burn operates on the claim, and the premium is not on the claim.

Now discount it honestly. The payoff is negatively correlated with the states the holder fears, which is the point; the *delivery* is positively correlated with them. $\Delta(R)$ spikes under energy

interdiction, hardware denial, network filtering, and jurisdictional pressure (Section 5.5†, Section 17.10.1†), the states that impair the protocol's ability to deliver at all. Red Line 13 reads that as a threat to the hinge; it is also a discount on the hedging premium. The holder is asked to pay for insurance whose underwriter is exposed to the insured event.

The slots are occupied, too. Gold already holds "works when the network does not" and Bitcoin "works when the state does not". A triad asset needs power, reachable networks, and reference hardware to deliver anything, so it hedges a narrower band: conditions adversarial enough to need proofs and privacy, not severe enough to take the grid. That band may still be the important one: repression that stops short of collapse is what Part I† argues is most likely. Sovereign optionality (Section 17.8†) exists to widen it, and disruption-adjusted VERIFYPRICE (Section 17.10.2†) to price whether it has.

Unsized. The mechanism is ordinary: service flows to a bearer are common in asset pricing. The **holder-side flow is nonetheless unsized**. Nothing in the thesis estimates the convenience yield a triad asset could command, and no honest estimate exists before there is a market to observe one in. That refusal is why the chain-strength rating for the store-of-value premium is Medium (conditional), and why the claim sits on link 9 rather than on fee accrual. The question is meant to be settled, not asserted: Section 26.8.6† publishes 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.

15.3 Denomination

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 unit-of-account question, the one monetary function the other instruments do not measure: not whether anyone holds the asset or pays with it, but whether anyone *chooses to quote* in it.

It is also the stickiest function. A store of value can be abandoned in an afternoon: the holder sells and is out. A unit of account is welded into contracts whose counterparties would all have to move at once, each eating the basis if it moves alone, which makes it a coordination device, hard to acquire and correspondingly hard to lose [Doepke & Schneider 2017].

The trap, first. A naive denomination share reads one hundred per cent on launch day and means nothing, because every triad service is priced in the base token by design (Section 9.3) and Condition 1 of the Value Capture Lemma requires it. Mandated denomination is a routing rule. Everything of interest is 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 is not settling. A contract quoted in dollars and settled in the asset is a payment-rail observation, and it is the configuration essentially all crypto commerce has taken. **Bitcoin is the most successful new store of value of the last century on the measures this thesis uses, and after seventeen years it denominates almost nothing.** Either denomination is not necessary for moneyness, or bitcoin's moneyness is incomplete along a measurable axis. The thesis takes the second reading and concedes it is contestable.

The reading rests on one asymmetry. What bitcoin buys is produced by an economy that prices in dollars, so a bitcoin quote imports basis risk against every input the seller pays for. In a triad protocol the unit and the goods come from *one* system: a prover paid in the unit, posting collateral in it and buying capacity priced in it holds a naturally two-sided position, exposed to basis only on the legs that reach outside. That is a coordination path to denomination a pure bearer asset lacks. It does not say the path will be taken.

The walled-garden test. Airline miles and in-game currencies denominate large internal economies and are not money, so what matters is where the counterparty stands. Tier I contracts have a protocol service on both legs, Tier II on one, Tier III on neither. Monetary content rises from I to III, the tiers are never blended, and Tier I is context rather than evidence: a protocol scoring highly there alone has built a company scrip with excellent telemetry. At the boundary the share is read against the dollar list price the buyer already has for compute.

The shape it predicts. Nothing prices in a volatile unit if it can avoid it, so denomination should appear **first at short tenor** and **first in two-sided positions**, where the variance nets, and last, or never, in long-dated one-sided contracts. A share rising uniformly across buckets

reads as subsidy or mandate. The cost that makes volatility prohibitive is the cost of carrying the exposure, which is what Section 15.1 lowers: **information-insensitivity lowers the carrying cost of denomination**.

Confirming, not leading. When populations abandon a domestic unit, savings dollarize years before prices do. If that ordering holds here, free-choice denomination confirms monetary adoption on the macro path rather than warning of it. It leads only on the protocol's own contract layer, where Tier I and Tier II contracts can be counted from Phase I onward. What is early is the measurement, not the phenomenon.

Instrument and falsifier. The share is published quarterly on the Value Capture Board (Section 26.9†) on the custody and reproducibility rules Red Line 14 sets, so the party collecting the fee may not classify the contracts, beside a quote/settle matrix, a tenor and sidedness split, a subsidy disclosure (a purchased quote is not a chosen one) and redenomination at renewal. Red Line 18 reads the result in two conditions, outsiders declining to reckon in the unit and participants declining on unmandated contracts, each with the consequence its box sets (Section 17.7).

The section does not claim denomination is necessary for moneyness, or likely; it raises no rating in Table 9.1 and promises no early warning on the macro path. It answers the question the other instruments cannot: what would be visible at Phase I that distinguishes this from every service token that ever shipped? 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 16

Stages, Phases and Macro States

The mechanisms say where a premium would come from. This chapter says how it would arrive, if it does: the combined thesis as separable stages from need to price, the negative case of useful infrastructure without money, the phases of adoption, and the policy states they unfold in.

16.1 From Need to Price

16.1.1 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 Part I† develops.

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.

Stage B is where the monetary object can fail while the system stays useful, and that failure has a recognizable shape.

16.1.2 The Negative Case: Useful Infrastructure Without Money

If the triad remains useful infrastructure but never becomes money, it looks like this:

- 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 stack can win without the coin ever becoming money.

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.

Which outcome arrives, the stack with money or the stack without it, is read along the adoption curve.

16.2 Four Phases of Adoption

The curve runs as four overlapping phases, each pulled in by a different constituency: cypherpunks, then allocators, then developers who stop noticing the stack, then states. They vary by geography and sector, and they are **metric-anchored, not calendar-anchored**: a phase has arrived when its telemetry shows it, not when a date does. Calendar horizons are omitted on purpose, since they age badly and invite cheap dunking. Each phase has an exit gate, the reading that opens the next, and a failure reading, the one that sends it back or lands on a red line.

Phase I: cypherpunk-led. The people who already live in the future (open-source cryptographers, hardware hackers, privacy-chain communities, early AI and ZK builders) prove that the loop of create, prove, settle, and verify works **without permission**. Their artifacts are reference implementations and receipts: open-admission prover markets publishing live VERIFYPRICE dashboards, testnets that measure rather than assume whether verification overhead stays inside its band on commodity hardware, public receipt ledgers, and BTC↔XMR swaps that people actually use, the one non-custodial privacy corridor deployed today. Most users are still speculators and hobbyists. A few things nonetheless become facts rather than hopes: proofs can be bought from permissionless markets, paid for over privacy rails without custody, and checked by anyone with a laptop.

Exit: VERIFYPRICE for ≥ 3 canonical workloads with ≥ 6 months of history and p95 inside its SLO under stress; VERIFYSETTLE for ≥ 1 corridor at $\geq 95\%$ success, with no protocol-

attributable refund failure and overall refund safety $\geq 99.5\%$, every failure root-caused; VerifyReach published for major regions with real multi-path reachability; and ≥ 1000 receipts a day. The question then stops being whether this can exist and becomes whether it can scale. *Failure:* verification cost above its SLO for three consecutive months without credible remediation is Red Line 1's reading, a refund failure the protocol caused on an admissible route is Red Line 2's, verification concentrated $> 70\%$ on one profile or jurisdiction for three consecutive months is Red Line 3's once Phase II has opened, and receipt data going dark leaves the thesis unverifiable.

Phase II: allocator-led. Fund managers, corporate treasuries, exchanges, and large platforms start to treat privacy and proofs the way they once treated bandwidth: as operating costs, not science projects. They **reclassify**. A bank's AI risk model runs on verifiable compute because regulators ask for receipts; a media platform attaches provenance because the liability of not doing so is too high; a trading venue settles privately rather than leak its order book. None of them has converted to cypherpunk ideology. Compliance and risk management now require math, not memos. **Budgets stop paying for promises and start paying for work that anyone can check.**

Exit: fee coverage $\geq 30\%$ of the security budget for 12–24 months, counted as retained fees with burns excluded; budgeted workloads outnumbering speculative ones in the receipt tags; VERIFYSETTLE inside its SLOs through at least one policy shock; capex and opex lines that name verification; and, where agents originate material demand, Principal Map readings inside their pre-declared bounds (Section 26.8.9†). The 30% sits roughly an order of magnitude above anything the closest precedents reached (Section 34.1.3†). It is set where the security budget stops being issuance-funded, a deliberately hard target with a known failure rate.

Allocators can arrive two ways that look alike on a price chart and nothing alike in the telemetry. Native adoption buys, stakes, and posts the base asset, pays native fees, and holds self-custodied claims; it moves every series on the Value Capture Board. Wrapper-led exposure buys funds, treasury-company equity, and futures; it may precede native adoption, seed it, or permanently replace it, and only VerifyFlow says which (Section 26.8†). Demand from software agents reads as native on the boards; whether it is monetary adoption depends on whose principal issues the delegations, since a fleet's demand is exposure by another route (Chapter 6). Assets under management, market capitalization, price, and wrapper launches therefore do not qualify the system for Phase III.

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.

Failure: fee coverage back below 30% with a workload mix $> 80\%$ speculative returns the system to Phase I, and a deeper, longer collapse is what Red Line 5 reads. The fee base is also late by nature. Budgeted verified compute is disproportionately a deployment-phase demand, one that plausibly becomes measurable only on the far side of a financial clearing event the thesis does not predict and cannot time (the capital-cycle clock of Section 16.3), so through a disorderly period the fee base should look worse before it looks better. That is a re-timing, not a relaxation: the threshold and the regression gate are unchanged, and Condition B of Red Line 14 reads the opposite case, a clearing event that came and went with no fee base behind it. **Phase III is not reachable on a 36-month horizon, and every deliverable that can be dated is a Phase I artifact.**

Agents can also invert the order. They need no treasury committee, only a receipt schema, a refund-safe corridor, and a unit anyone can check, which is Phase III's defining property shipped by construction. No gate is relaxed for that, and composability that is one operator's SDK is a fleet, not a phase.

Phase III: composability-led. The stack disappears into infrastructure, the way TLS did on the web. New applications emit claims and receive receipts because it is easier than hand-rolling trust, frameworks ship verification modules, and payments that cross borders route quietly over privacy rails with receipts that satisfy auditors but not censors. From the outside nothing dramatic happens; what changes is the **default**. Who carries that default to the median user, a carrier's channel or the user's own agent, is read as the Cross-Perimeter Share at the consumer edge (Section 6.2), a series not yet computable. This phase can show **durable cash-flow accrual** to the asset. Whether a monetary premium sits on top is a separate question, decided by the holder-side flow (Section 15.2) and not answered by price in either direction (Corollary 13.1).

Exit: the base asset as collateral or reserve in ≥ 3 independent native venues, with the Native Monetary Buyer Map showing a self-custodied, loss-bearing cohort rather than inferring one from lockup; VERIFYPRICE and VERIFYSETTLE stable through macro stress; a cohort holding the base asset for years, with Work Credit inventory reported separately as service demand; and fee coverage $\geq 50\%$, sustained. The base asset has then advanced as a monetary candidate, and Work Credits have not changed instrument class. *Failure:* the Value Capture Board reads Phase III on one operator's traffic.

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 IV: policy-led. Once Phase III holds through macro stress, the curve turns political. Some jurisdictions codify cryptographic receipts as evidence, mandate provenance for certain media and AI systems, and recognize privacy rails with viewing keys as compliant infrastructure. Others keep their chokepoints. If the rails stay non-custodial and censorship-resistant, and proof and compute markets stay globally addressable, capital and talent can route around them the way data routed around telcos. *Real once:* receipts codified in ≥ 2 jurisdictions, regulation that references proof primitives rather than records, public institutions publishing their own receipt and corridor metrics, and volume migrating measurably toward jurisdictions that respect lawful privacy. In the best case each wave of repression pushes savings onto rails where verification is cheap, and each wave of adoption deepens those markets. That is a best case, not a forecast, and the premium in it remains the unsized quantity of Section 15.2.

Any phase can fail backward. A slide from Phase II back to Phase I is a failure of adoption, not of the technology, and the telemetry exists so the difference can be seen. Nor is the curve guaranteed. Verification cost can creep, hardware or router oligopolies can re-centralize control, regulation can pinch the ramps harder than expected. If the stack stays anchored in verifiable machines, keeps verification cheap and public, treats privacy as infrastructure and forces its own claims through telemetry and receipts, this is the curve it can plausibly aim at.

16.3 Macro Policy States

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, so the model names states and observables rather than a dated sequence. Its market-repair distinction and repair sequence are adapted from Michael Green's commentary [Green 2026c] [Green 2026d], which supplies mechanisms and scenarios, not adopted quantities or tactical forecasts.

State 1: market repair succeeds. Buybacks, maturity changes, regulatory repair, or lower front-end rates restore market function while long-dated inflation compensation stays contained. 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, power, shipping, or industrial bottlenecks lift long-dated inflation compensation. The financial system may need easier policy while the physical system argues for restraint, and Layer 0 costs rise just as demand for neutral assets may. Long-dated breakevens and inflation swaps are the tripwire.

State 3: repair fails and repression follows. No durable marginal buyer returns, and policy shifts toward captive demand, preferential regulation, negative real returns, capital controls, or restricted exit.

State 4: competent closed-stack stabilization. The state coordinates energy, industry, compute, identity, payments, and duration warehousing; functionality and collateral stability improve while user custody and exit contract. This is not collapse. It is the open stack's strongest competitor.

Transitions are conditional and read from observables, and none of them supplies a tactical forecast for bonds, currencies, or commodities. Three readings cut across all four states. **Household transmission:** a state can be stable at the sovereign level while its households deplete, and a State 1 with exhausted buffers moves toward whatever regime the depleted median will vote for (Section 3.3). **The capital cycle** is a clock rather than a state [Perez 2002], borrowed for one mechanism and not for its historical periodization, which the thesis does not adopt: installation finance outruns what the economy can absorb, a Turning Point reprices the installed base, and a Turning Point arriving in State 2 meets a central bank with less room to absorb it than one arriving in State 1; the phase gates of Section 16.2 are read against it. **Duration absorption**, not the debt ratio, decides whether repair holds, because the marginal buyer is increasingly demand commissioned by benchmark rules and financed in repo (Section 3.1.3). None of this times the transition; it says the plumbing series would show it first (Appendix I.6†), and the long yield is not among them.

Recession first, repression later is a path, not a promise.

Part VII
The Evidence

Every link in the chain is a conditional, and a conditional is worth stating only if a reading can break it. This Part sets out those readings: where the stack is exposed, which readings would tell against the thesis, the lab that would take the measurements, and what the bet comes to.

Chapter 17

Risk Analysis & Failure Modes

The risks to this emerging triad are not merely technical; they are structural, political, and economic.

17.1 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.

17.2 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.

17.3 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.

17.4 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.

17.4.1 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.

17.5 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.

Exposure is not yet failure.

17.6 Eighteen Ways to Be Wrong

The thesis is falsifiable, and it says in advance where. Each red line names a measurable condition and what follows if the condition persists. The consequence carries a severity label, and the label is not a tone choice. A line that **kills** falsifies the thesis outright: the hinge of the chain is severed. A line that **weakens** removes one leg of support, and the thesis survives only in reduced form, to be restated.

A single bad reading is an incident. A breach that persists without credible remediation is

Red line	What it reads	Consequence
RL1	*Verification Affordability Breaks.* The cost of checking core workloads, held against its SLO	Kills
RL2	*Refund Safety Breach.* Whether an honest party always leaves a swap whole, and a broken corridor is delisted	Kills
RL3	*Verification Monoculture.* How much verification runs on one vendor, hardware profile or jurisdiction. Reads only after Phase II entry	Kills
RL4	*Telemetry Capture.* Whether the receipt data stays public, checkable and out of any one party's hands	Kills
RL5	*Fee Coverage Collapse.* Retained fees against the security budget, and how speculative the workload has become	Kills
RL6	*Value Capture Failure.* Whether native fees, burns and lockups grow with triad usage, or bypass channels take it	Kills
RL7	*Legal Incompatibility.* Whether lawful users in major jurisdictions can use the privacy rails, with lawful-privacy patterns adopted	Weakens
RL8	*Governance Capture.* Hard limits on governance, the concentration of voting weight, and how routine the emergency path becomes	Kills
RL9	*Wrapper Dominance Becomes Monetary Substitution.* Wrapped and custodial exposure outgrowing native fees, settlement and lockups. Reads only after Phase II entry	Weakens
RL10	*Physical Infrastructure Opacity.* Whether claims about facilities can be checked, and how cheaply. Reads only after Phase II entry, once a baseline exists	Kills
RL11	*AI Enclosure.* The share of verified compute capacity held by the largest hyperscalers, closed TEEs or one jurisdiction. Reads only after Phase II entry	Kills
RL12	*Agency Failure.* Forced disclosure in everyday onboarding, and the non-custodial share of native volume; a watch clause on agent-held balances trips nothing yet	Kills
RL13	*Energy Sovereignty Failure.* Network sovereign optionality against its published floor, and how much capacity sits in host states that curtail or ration power against verification workloads. Reads only after Phase II entry	Kills
RL14	*The Capturable Wedge Closes.* Whether buyers pay for verification at all: the wedge against the take rate (A), fee turnover against throughput (B). Either breach alone trips it	A: kills; B: kills
RL15	*Native Collateral–Capacity Spiral.* A drawdown that drains collateral, operators and deliverable capacity at once, with service not recovering. Reads only after Phase II entry	Kills
RL16	*The Convenience-Yield Null.* Whether holders charge more to part with units in the corridor than in benign states. Reads only at Phase III scale, with a deep, continuously quoted market	Kills only if powered
RL17	*Information-Sensitivity Reversion.* Whether independent lenders disagree less about the asset's haircuts than about an opaque twin's. Reads only with a panel of at least five independent lenders	Kills
RL18	*The Denomination Null.* Whether outsiders (A) and insiders (B) reckon in the unit, not merely settle in it. Reads only above a free-choice volume floor, and B has a second floor of its own	A: weakens; B: kills

Table 17.1: The eighteen red lines, each consequence as its box gives it. Before a line starts to read, its reading is published and trips nothing.

a verdict, and the response to it is fixed before it happens. The breach is declared an incident, and a remediation plan is due within 14 days; if remediation has not worked after 90 days, the public boards reclassify the asset from "SoV candidate" to "speculative/experimental" on a line that kills, and restate the claim without the removed leg on a line that weakens. Either is a published reading, not anyone's announcement. Several lines breached at once open crisis mode, with an independent review whose results are published.

The first fifteen lines read principally on links 2 through 8 of the chain (Section 9.1): the service path, the capacity, the fee accrual. The last three read on link 9, the monetary mechanisms themselves (the holder-side flow, pledgeability and denomination), and this edition prints their boxes in full in Section 17.7. A falsification regime with nothing aimed at that link would leave its sharpest criticism unanswered. These three close the gap rather than argue it away.

The table compresses and the boxes govern. Thresholds, clocks, instruments and reading notes are in Section 30.6† of the full edition. A nineteenth line, on the concentration of operating rules among native holders, is published there as a candidate: it is not one of the eighteen, and it trips nothing in this version.

These are not punishments. They are truth in advertising: an asset that claims to store value must show that it does, or stop claiming.

17.7 The Three Monetary Red Lines

The red lines on the three monetary mechanisms follow word for word, as the full edition states them.

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 16.1.1: 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 lo-

cating 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 16.3), 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 7.1 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 9.2) 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 13.2.7 shows that fee accrual is not a source of monetary premium, and Section 15.2 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 7.1). 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 15.1) 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.

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 15.1). 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 15.1 is absent, and cheap verification has bought operational convenience without buying moneyness. This red line reads on link 9 of Section 9.1, which the fee, capacity, and coverage lines do not.

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 15.1 and Section 15.2

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 16.2), $\geq 25\%$ 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 15.3), 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 exception. 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.

Chapter 18

Kardashev Labs: Build and Measure

The red lines say what to measure; someone still has to build the stack and take the readings. This chapter is for the reader deciding whether to fund that work rather than whether to believe the argument.

18.1 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 7.4.1). 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 16.2). And the residual monetary premium (the three mechanisms of Section 15.1, Section 15.2, and Section 15.3) 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.

18.2 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 15.1). 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 7.1). 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 14). 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.

18.3 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 (Section 32.3†), 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 (Section 32.5†) 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 19

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 18): 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 18.1).

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.

19.1 A Quiet Arrival

If it does work, the moment of success will not look like an ICO or a conquest. It will arrive the way administrative repression does: as paperwork.

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.

19.2 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 more than a higher price: 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.

19.3 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 7.4.1). 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**.

Does the premium follow? The thesis does not know; the lab exists to find out. 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.

19.4 Equip, Don't Manage

The purpose of the stack is to preserve the capacity to act when institutions become coercive, synthetic, or brittle, not to make people more legible to them. 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 only the third still under test.

This closes the circle opened in Section 3.3. 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 6.6). 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.*

[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: Communications 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.

[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 16.3 (the capital cycle as a cross-cutting clock on the four states), Section 29.1† (the window closing by event rather than accrual), Section 13.2.2 and Section 11.1† (financial repricing as a deflation channel distinct from technological deflation), Section 16.2 (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 15.1 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 15.1 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 15.1 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 14 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 14 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 15.3 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 15.1 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/1808571. *Cited in Section 15.1 and Section 13.2.6 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; at the fork it is applied to the open stack as a credible entrant that can discipline incumbents' rents without taking share. 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.*

[Teece 1986] **Teece, D. J.** (1986). Profiting from Technological Innovation: Implications for Integration, Collaboration, Licensing and Public Policy. *Research Policy*, 15(6), 285–305. https://doi.org/10.1016/0048-7333(86)90027-2. *Cited in Section 13.2.6 and Section 7.3 for one mechanism only, stated in its abstract: that when imitation is easy, the profits from an innovation may accrue to the owners of complementary assets rather than to the developers of the intellectual property. Used to name who earns the return on an open protocol: its own coordination assets at the fork, and the custody and distribution that wrappers own. Read: the abstract and the opening section (pp. 285–286). No result about integration, licensing, appropriability regimes, dominant designs, or trade policy is imported, and no threshold in this thesis is derived from it.*

[Coase 1937] **Coase, R. H.** (1937). The Nature of the Firm. *Economica*, 4(16), 386–405. https://doi.org/10.1111/j.1468-0335.1937.tb00002.x. *Cited in Section 6.3 for one mechanism only: that the boundary of the firm sits where organizing a transaction internally costs as much as using the market. Verification cost is read as one such transaction cost. No other result is imported, and no threshold in this thesis is derived from it.*

[Barzel 1982] **Barzel, Y.** (1982). Measurement Cost and the Organization of Markets. *Journal of Law and Economics*, 25(1), 27–48. https://doi.org/10.1086/467005. *Cited in Section 6.3 for one mechanism only: that when the attributes of what is traded are costly to measure, markets grow arrangements to economize on measuring them, such as warranties, packaging, product ratings, and certification, which would be superfluous if product information were costless (pp. 27–28). $r(W)$ is read as that measurement cost for machine work. No result about any particular market is imported.*

[Manidis & Roberts 2026] **Manidis, W., & Roberts, J.** (2026, October 29). *The Commanding Heights.* Web edition. https://commanding-heights.com. *Cited as a position: in Section 7.4.3 for the case for owning a durable toll, a scarce, controllable position whose revenue grows with the nominal value of the activity passing through it; in Chapter 18, where the base asset is asked whether it is one; and in Section 16.2 for the account of diffusion in which the installed base is the position worth holding. No quantity or forecast is taken 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 15.2 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 15.2 is to the mechanism, not to the setting.*

Market Structure, Flows & Inelastic Markets

[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.5†.*

[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 16.3. 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 16.3 and the exhaustion gradient of Section 3.3); 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 7.4.4). 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 15.1 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 15.1 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 17).* ***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.5†: 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.*

[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 7.2.1 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.*

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 8 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.

[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 8.1). 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.5†.* https://fftt-llc.com/.

[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 8.1 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.*

Note on the synthesis. The layered reading in Chapter 8—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.

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/notes/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.*

[Brynjolfsson et al. 2026] **Brynjolfsson, E., Chandar, B., & Chen, R.** (2026, August 12; first circulated 2025). Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence. Working paper, Stanford Digital Economy Lab. https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine. *Cited in Section 7.4.4 for the direction of one finding, as stated in the abstract of the revision dated 12 August 2026: in ADP payroll data through June 2026, employment of workers aged 22–25 in AI-exposed occupations stands below where it would otherwise be. The authors describe the results as early and descriptive rather than causal. The headline figure has changed between versions, and none is used here. Working paper, not peer-reviewed.*

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 3.2 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 3.2 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 3.2 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.

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.*

[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.*

[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 13.2 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."*

Privacy Rails: Zcash, BTC↔XMR Swaps

[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/GHSA-ghc3-g8w4-whf9. *Cited in Section 17.1 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.

Proof Markets / PoUW / Privacy-First L1s

[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 16.2, 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.*

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.

Censorship Measurement & Circumvention (Layer 1 Prior Art)

[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."*

[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.*

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.