

AfterFiat

The Load-Bearing Thesis

*A concise, conditional monetary thesis
for verifiable digital infrastructure*

Jason St George

Privacy, Proofs, and Compute may support an agency-preserving bearer asset. Whether they do—and whether they should share one asset—is an empirical question. The thesis does not ask to be believed; it asks to be measured.

Version 3.5

13 October 2026

Status of This Edition

This is an author-maintained condensation of the canonical full AfterFiat version 3.5 thesis. It is designed to stand alone and preserves the full edition's central mechanism, qualifications, hierarchy of instruments, objections, falsifiers, and experimental program. It is not a replacement for the research dossier. **If wording, scope, definitions, thresholds, or interpretation differ, the canonical full version 3.5 thesis governs.**

© 2026 Jason St George. Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0). "AfterFiat," "VerifyPrice," "VerifyReach," "VerifySettle," "VerifyFlow," and related marks remain trademarks/service marks of the author and are not licensed under CC BY 4.0.

Citation: Jason St George. *AfterFiat: The Load-Bearing Thesis*. Version 3.5, 13 October 2026.

Contents

1	The Claim Under Test	1
2	Threat Model: What Must Survive	2
2.1	Adversary classes	2
2.2	The competent closed-stack competitor	3
2.3	Standing inversion	3
3	Three Monetary Mechanisms, Each With a Falsifier	3
3.1	Pledgeability without due diligence	5
3.2	Denomination	5
4	Value Capture Is Not Monetary Premium	5
4.1	Protocol value and market price are different machines	6
5	The Eleven Premises and Nine Links	7
5.1	Eleven premises	7
5.2	Nine-link conditional chain	7
6	Topological Scarcity, Delivered Capacity, and the Corridor	8
6.1	Delivered Verified Capacity	8
6.2	The Pressure–Capacity Corridor	9
7	Native Buyer Quality: Demand Is Not an Anchor	9
8	Triad Coherence and Instrument Hierarchy	10
8.1	Triad Coherence Test	10
8.2	The hierarchy is constitutional	11
9	Strongest Objections and What Is Not Proved	11
9.1	The objections that remain load-bearing	11
9.2	What this edition does not prove	13
10	Eighteen Red Lines	13
11	Four-State Macro Model	15
12	Minimum Viable Experimental Roadmap	16
12.1	1. Receipts and verification harness	16
12.2	2. One privacy corridor	16
12.3	3. One proof factory and canonical starter set	17
12.4	4. DVC and physical evidence	17
12.5	5. Developer surfaces and neutral routing	17
12.6	6. Monetary and collateral experiments	17

12.7 7. Useful-work pilots last	17
12.8 Decision rule	18
13 What Is Being Bought: Stack First, Coin Optional	18
14 Conclusion: Equip, Do Not Manage	18

1 The Claim Under Test

AfterFiat begins from a narrow observation. Modern money, communication, identity, and compute increasingly depend on soft guarantees supplied by institutions whose rules can change: a custodian will honor 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. Debt arithmetic, synthetic media, administrative integration, and concentrated compute make those guarantees more valuable and more conditional at the same time.

The response is not a forecast of fiat collapse and not a claim that useful services naturally become money. It is a test:

Under sustained administrative and financial repression, a bearer base asset may earn monetary premium if holding it preserves private settlement, portable proof, and access to verified compute after ordinary legal, custodial, and platform 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 duration-bearing infrastructure credit must remain separate. Privacy, Proofs, and Compute may share that asset, but whether they should is an empirical architectural question rather than a premise.

Every important verb in that statement is conditional, and one condition comes first. The infrastructure can succeed without any new money; the question is whether possessing this particular bearer asset gives an advantage the infrastructure alone does not — above all, whether a balance held in advance beats holding Bitcoin and converting on the day. The thesis investigates that rather than assuming it. The relevant base reality is the networked ability to deliver three distinct services under adversarial conditions:

- **Privacy:** non-custodial settlement with privacy by default and user-controlled, scoped disclosure.
- **Proofs:** portable attestations of computation, provenance, and policy predicates that remain cheap for an independent party to check.
- **Compute:** useful machine work wrapped in guarantees whose verification cost stays well below production cost.

These services are not themselves money. Nor does the thesis promote every object associated with them into a monetary instrument. The **base asset** is only a conditional monetary candidate. **Work Credits** are typed, prepaid claims on service reserved with named providers. **Receipts** are evidence. **Project notes** are explicit duration-bearing credit. **LP and staking positions** are derivatives or operating claims.

Gold, silver, Bitcoin, and physical hard assets remain load-bearing bridge assets while the proposed stack is immature. Bitcoin's work function is monetarily superior to proof-of-useful-work in one decisive respect: nobody outside Bitcoin buys the hash, so no external buyer can withdraw, mandate, or subsidize its demand. Proof-of-useful-work introduces a buyer and therefore trades monetary objectivity for capacity relevance. AfterFiat supplies different goods; it does not establish that they will carry a larger monetary premium than the assets already occupying adjacent roles.

2 Threat Model: What Must Survive

The protected object is not a token balance in isolation. It is the complete *Create/Compute* $\rightarrow$ *Prove* $\rightarrow$ *Settle* $\rightarrow$ *Verify* loop and the physical and institutional substrate beneath it. Four properties must survive:

1. **Integrity of state and receipts.** Computation, provenance, settlement, and consensus artifacts must not be silently corrupted.
2. **Confidentiality of flows.** Privacy must survive ordinary use; optional disclosure must be scoped by the holder—per payment, per epoch—and cannot mean a universal transaction graph.
3. **Availability and neutrality.** Users and providers require open admission, reachable routes, non-seizable settlement, and practical exit.
4. **Verification economics.** Checking a claim must remain cheap in absolute terms and cheap relative to producing it. Otherwise public verification degrades into a trusted-verifier priesthood. The full edition distinguishes three *modalities*: succinct proofs (M1), whose checking cost is independent of the work; algebraic or probabilistic checks (M2), such as Freivalds’ matrix-product test; and replication or attestation (M3), which rests on a committee. Only M1 and M2 carry the monetary mechanism; M3 supports service delivery and nothing more.

2.1 Adversary classes

The **rational sovereign under stress** uses the levers it controls: negative real rates, capital controls, custody rules, identity mandates, hardware standards, app stores, telecoms, clouds, and interconnection queues. The likely form is often not prohibition but **administrative repression**: formally optional rails become practically mandatory through defaults, compliance rules, tax treatment, benefits, procurement, platform terms, and institutional mandates.

Platforms and strategic patrons can centralize provenance, decide which proofs count, privilege national champions, and convert private compute into protected quasi-public infrastructure. This *enclosure by rescue* may be more dangerous than overt attack because the closed competitor can be well funded, technically excellent, energy-privileged, and genuinely useful.

Economic and compositional adversaries include prover cartels, router capture, liquidity games, dealer hedging, wrapper redemptions, leverage, and common operating rules. No participant need be malicious. Locally rational behavior can still produce correlated sales, recursive leverage, procyclicality, concentration, and prices that reveal little about native use.

Hardware, supply-chain, and physical adversaries exploit biased randomness, closed attestation roots, firmware control, export restrictions, curtailment, discriminatory tariffs, cooling constraints, transformer queues, network cuts, or load priority. A state that finds verification awkward to ban can decline to energize proving or make unprivileged verifier hardware unobtainable.

Cryptanalytic, governance, and vaporware risks remain ordinary but fatal: proof-system breaks, captured update keys, unverifiable benchmarks, silent parameter changes, and performance claims without reproducible data.

2.2 The competent closed-stack competitor

The real competitor is not necessarily a failing state. It is a competent closed sovereign stack integrating energy, industry, compute, identity, payments, and duration warehousing. Such a system may win on cost, latency, uptime, and construction speed. AfterFiat's claim is not capability superiority. It is that non-custodial settlement, private-by-default disclosure, portable identity, independently checkable receipts, and practical exit are distinct properties. A closed stack withholds them because that withholding is what makes it closed.

The scoreboard must therefore include agency and exit. Two systems can show identical throughput and uptime while giving opposite answers to the question: can the participant leave with keys, assets, proofs, and usable service?

2.3 Standing inversion

The marginal buyer of cheap verification is increasingly a machine delegating to another machine, several hops from a principal who may not have specified the task. That buyer brings a failure mode the stack itself creates. Agents can acquire wallets, receipt-built reputations, and collateral positions—every measurable property of a good holder—while human title over them stops binding: the key root is not held by a person, the delegation ceiling can be raised by the agent or its operator, and revocation is a request rather than a protocol operation.

There is a second route to the same place on which every key stays with a person and is never used. A bad outcome after following an agent's standing instruction is charged to the instruction; a bad outcome after overriding it is charged to the person; and a stack that makes every decision cheap to reconstruct widens the gap. Title stays intact while judgment is delegated with the keys. Every existing board reads green on both routes. The full edition therefore classifies native demand and holders by ultimate principal and reads standing as a pair: the share of agent-held balances a person can still revoke, against the policy concentration of the same cohort. Humans holding every root while the holdings move as one policy is standing exercised by nobody.

3 Three Monetary Mechanisms, Each With a Falsifier

Monetary premium is easily defined as a residual—value in excess of discounted cash flows—and fees, burns, and collateral correctly fail to reach it. A residual admits no affirmative evidence. The claim therefore rests on three positive mechanisms, each carrying an instrument and a red line that retires it. The first is developed below; the other two are stated after it.

Let R denote regime pressure and let $\Delta(R)$ denote the value of the protocol's differential properties over the best institutional substitute. In a benign state, a solvent, suable indemnitor is an excellent substitute for a cryptographic proof. Courts function; custodians redeem; platforms serve the user; hyperscalers offer low cost and contractual recourse. For the median commercial workload, $\Delta(R)$ is therefore probably thin.

The differential changes when the substitute cannot be sued, will not remain solvent in the relevant state, is itself the adversary, or causes a loss that damages cannot repair—censorship, seizure, deplatforming, or disclosure. Then the ability to transact, prove, hold, and exit can become valuable

precisely because the institutional substitute has weakened.

State contingency creates two mechanisms that must not be confused:

Countercyclical fee cash flow.

If demand and fees rise in bad states, the fee stream can deserve a lower risk-adjusted discount rate. That is real value, but it remains inside a discounted-cash-flow valuation. A countercyclical utility is a better utility; it is not thereby money.

Holder-side service flow.

A bearer who already holds the base asset may retain the practical ability to transact, prove, hold, and exit when substitutes fail. This service accrues by virtue of holding rather than as a distribution. It is analogous only in mechanism to a convenience yield: the holder receives an option-like service that is absent from a claim on future delivery.

Regime-Contingent Convenience Yield

The first mechanism: the holder-side service generated by the state contingency of $\Delta(R)$ —the bearer’s practical ability to transact, prove, hold, and exit when substitutes weaken— plus credible non-discretion in issuance and bearer holdability. It is not the countercyclicity of fee cash flows and is not generated by burns, collateral, or fee share.

This establishes a plausible mechanism, not a magnitude. The holder-side service is unsized. No honest estimate is available before a market and repeated pressure episodes exist. Its delivery also has a defect: demand for the hedge may rise in the same states that damage power, hardware, networks, liquidity, and exit. The asset is insurance whose underwriter is exposed to the insured event. That defect is why deliverability is part of the monetary mechanism rather than an engineering appendix.

The unsized claim is nonetheless measurable, and the full edition specifies the instrument: base-asset lending rates (the price of temporary access to bearability), forward and perpetual basis, and wrapper basis, published against the regime-pressure index and read as the increment in the pressure corridor the thesis predicts — an inverted U, not a slope. Red Line 16 kills when the upper bound on that increment lies below a pre-registered minimum, and only when a pre-registered power calculation shows it could have seen the effect. Unsized with no instrument would be an article of faith; unsized with a published instrument is an open empirical question.

Why the hedge bid reaches the base asset and not Work Credits. A holder hedging loss of access could instead hold Work Credits for the service and Bitcoin for exit. Work Credits are typed, workload-specific, retire on redemption, and are claims on the *specific* capacity that fails in exactly the states where the hedge is needed; Bitcoin hedges exit but not access to verified compute or proofs. The base asset is the only instrument that is bearer, workload-agnostic, and the unit in which Work Credits are priced. If holders nonetheless prefer Bitcoin plus Work Credits, or Bitcoin converted on the day, that is the outcome the access experiment and Red Lines 16 and 18 detect, and the thesis treats it as its default null rather than as an anomaly.

Why hold rather than convert on the day. That answers Work Credits and Bitcoin as an exit asset, and leaves the strongest alternative: hold Bitcoin and buy the base asset when the need arrives.

Against it, holding buys one thing — a balance that stays spendable where acquiring one is unavailable, too slow, or incompatible with the transaction's privacy and control terms, while native settlement and the service path still work. It buys no capacity, price, or priority; those are sold through Work Credits. The counterfactual is the best feasible Bitcoin-funded strategy, provider-side conversion included, and an access experiment decides it, run first as an analog on an existing network whose fees must be paid in its own asset. If Bitcoin-funded procurement reproduces the service within the same constraints without material disadvantage, the mechanism is absent for that workload and state. A pass is not yet a reserve: agent wallets and gas inventories are operating buffers, not an anchor.

3.1 Pledgeability without due diligence

The second mechanism is the property the collateral literature calls *moneyness*: an asset is money-like when no counterparty gains anything by investigating the particular unit tendered. Cheap public verification supplies that property for *unit quality*—whether this receipt, this pledged position, this capacity claim is what it purports to be—and leaves *payoff* uncertainty untouched. Bitcoin is the control case: its units have been perfectly verifiable since 2009, and its collateral haircuts nonetheless sit at 30–50% because they price volatility, which no verifier can compress. The prediction is therefore narrow. Controlling for volatility and depth, a position whose backing is publicly verifiable should show lower *dispersion* of haircuts across lenders and faster convergence after a shock than an opaque comparable—not lower haircut levels—among lenders who underwrite independently; agreement produced by a shared model does not count. Red Line 17 tests exactly that, with a lender-panel maturity gate, and is specified so that Bitcoin's own record (native units against wrapped) is a case it can pass or fail rather than one it fails by construction. Pledgeability attaches to the over-collateralized, capacity-backed *position*, not to the bare token, which is the first monetary—rather than engineering—reason Layer 0 matters.

3.2 Denomination

The third mechanism is the stickiest monetary function and the one Bitcoin has not acquired: whether anyone reckons in the unit on contracts nobody required them to write in it. Measured as free-choice denomination share, tiered by whether the counterparty sits inside or outside the protocol. Outsiders declining to denominate *weakens* the claim; protocol insiders declining to denominate their own short-tenor obligations—prices fixed for 90 days or less, held to a 25% floor at Phase III entry—*kills* it (Red Line 18). Long-tenor contracts, which the theory says move last, are published without a fatal threshold. Bitcoin, whose miners quote hashprice in dollars and pay their bills in dollars after seventeen years, would trip that condition on any reading now available and is the standing counterexample.

4 Value Capture Is Not Monetary Premium

The protocol can route service demand into a base asset through native fees, retirement of supply, operator collateral, and disciplined issuance. These mechanisms matter, but they answer where economic value goes, not whether the asset is money.

Value Capture Channels

In the reference design, service demand accrues to the base asset through five channels, and the design relies on all five holding together:

1. the asset is required for core fees at equivalent service quality;
2. a meaningful share of fees is burned or permanently retired;
3. capacity providers must post the asset as collateral;
4. gross issuance follows a schedule fixed ex ante, not governance fiat; and
5. users cannot obtain equivalent service through a bypass channel that avoids the asset.

Each is necessary to this design's route to accrual and none to value capture in general: a buyback funded by fees in another unit, pro-rata distribution, or native fees without native collateral capture value by other routes. None is sufficient for monetary premium.

A competitive service market can still support a transaction fee: the fee is a wedge on turnover rather than a claim on operator profit. But the sustainable fee is bounded by what the protocol is differentially worth. If c_p is protocol delivery cost, c_b the best bypass cost, f the fee, and Δ the buyer's willingness to pay for checkability, neutrality, non-discretion, and exit, routing requires

$$c_p + f - \Delta \leq c_b, \quad f \leq \Delta - (c_p - c_b).$$

The protocol begins with a cost handicap because proving and redundancy consume real resources. Bypass is therefore a magnitude, not a switch. The commercial question is whether the high- Δ tail—cross-jurisdictional, pseudonymous, politically exposed, or audit-critical users—is large enough to support a fee after the verification premium is paid. That tail is two populations, not one: supervised counterparties who need evidence, and unsuable counterparties who need absence, and each one's requirement is the other's disqualifier through the shared anonymity set. Revenue is not additive across them. A fork runs the same technology, so against a fork the bound is coordination value alone.

Even when the five conditions hold, they produce two familiar valuation objects: a claim on the fee stream and a balance-sheet lockup of required collateral. A burn is economically a buyback. Collateral is a lockup that holds while operators stay solvent and choose to participate; it places no bid, and it is not a premium. Neither mechanism makes Visa, a pipeline, or a clearinghouse into money, and neither does so here.

4.1 Protocol value and market price are different machines

Native value capture is service use $\rightarrow$ fees $\rightarrow$ burns, collateral, and operator demand. **Price capture** is demand for exposure arriving through funds, custodians, treasury companies, futures, options, swaps, leverage, passive allocation, or dealer hedging. Either can move without the other. Therefore

price appreciation $\nRightarrow$ monetary adoption,
monetary adoption $\nRightarrow$ immediate price appreciation.

VERIFYFLOW interprets price formation; it does not replace VERIFYPRICE, VERIFYREACH, VERIFYSETTLE, capacity telemetry, or native-holder measurement. Price is an output to explain, never evidence to cite by itself.

5 The Eleven Premises and Nine Links

5.1 Eleven premises

- P1. **Soft guarantees are weakening.** Money, media, identity, and infrastructure rely increasingly on conditional institutions, platforms, vendors, and compliance systems.
- P2. **Digital civilization has three unavoidable needs.** Private settlement, portable attestation, and verified compute are distinct scarce capacities.
- P3. **Those needs can become verifiable commodities.** Standard workloads, portable receipts, and cheap public checking can replace selected platform promises with tradable facts.
- P4. **Gross capacity is not deliverable service.** Installed power and hardware matter only through the complete surviving path to usable, settled service.
- P5. **A store of value requires more than utility.** Indispensability does not determine which asset, operator, or customer captures value.
- P6. **Value capture requires enforceable design.** Required fees, supply retirement, collateral, issuance discipline, and a non-trivial capturable wedge must be visible.
- P7. **Gross native demand is not a monetary anchor.** Fees, burns, operator balances, collateral, and wrappers do not prove a durable holder capable of bearing loss.
- P8. **The system must remain falsifiable.** Verification, reachability, settlement, capacity, economics, holder quality, and agency require public telemetry.
- P9. **Market price is not proof of adoption.** External financial machinery can amplify, suppress, or counterfeit the appearance of native monetary demand.
- P10. **Duration-neutral money is not duration finance.** The base asset must not become a coupon; project construction still requires explicit credit and a loss-bearing holder of time.
- 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 reliability mechanism that fits inside the hop. That demand is structural only where delegation crosses trust perimeters; its coercibility depends on who operates the agent; and whether a person still holds standing over it is a question the stack must answer rather than assume.

5.2 Nine-link conditional chain

- 1. **Persistent utility demand** exists for private settlement, portable proofs, and verified compute. It fails if the demand is cyclical or discretionary rather than structural.
- 2. **Demand standardizes into canonical work** with workload registries, hardware profiles, and service tiers. It fails if every workload remains bespoke.
- 3. **Work produces cheap-to-check receipts** while verification remains far cheaper than production. It fails when “anyone can verify” becomes “trust the prover.”
- 4. **Receipts represent stress-deliverable service.** It fails when nameplate capacity survives on pa-

per but substitution latency or common dependencies remove usable output.

5. **Receipts enable open markets** for proofs, routing, escrow, and capacity. It fails when provers, routers, or matching engines replace work with rent.
6. **Markets produce recurring fee flows.** It fails when fees depend on subsidy or speculative workloads rather than budgets.
7. **Fee flows accrue to a scarce base asset** under the five value-capture conditions and a positive capturable wedge. It fails when equivalent service bypasses the asset.
8. **A monetary-risk warehouse forms.** A persistent constituency retains self-custodied balances, bears losses, and adds through stress (spending on the service during stress counts as use, not exit). It fails when demand is acquired just in time, immediately sold, leveraged, wrapped, or procyclical.
9. **The anchored asset may earn a store-of-value premium** only if it remains liquid, neutral, verifiable, legally holdable, and agency-preserving, and if a held balance beats Bitcoin converted on the day. Passing the link creates a candidate; it does not size the premium or forecast price.

The weakest links are intentionally visible. Standardized useful work and proof/compute market formation remain immature; stress-deliverable service requires common-cause and substitution data not yet demonstrated at scale; non-bypassability is an unestimated magnitude; the native holder anchor is unproven; and the final premium remains medium and conditional.

6 Topological Scarcity, Delivered Capacity, and the 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 coexist. Crude is not diesel when refining is the minimum cut. Savings are not duration absorption when no balance sheet can own the residual risk. Installed accelerators are not verified service when power, cooling, firmware, network access, verification, or settlement fails. The transformation path, location, actor, and required horizon determine the usable service.

This is a mechanism-level synthesis. Physical-conversion and marginal-buyer examples help identify the topology; they are not treated here as facts about a particular geopolitical outcome or as tactical forecasts.

6.1 Delivered Verified Capacity

For canonical workload W , interval t , and disruption scenario s , define a directed graph $G_{W,t,s}$ from available energy and fuel to usable settled service. Every edge has surviving capacity $c_e(W, t, s)$ in a workload-specific service unit. Then

$$\text{DVC}(W, t, s) = \text{MaxFlow}(G_{W,t,s}, c(W, t, s)).$$

The constraints that bind at the optimum identify the active bottleneck. With gain factors along the chain, or several workloads sharing edges, the textbook max-flow / min-cut equality fails in general, so the report names binding constraints rather than a single cut. The unit is not a generic FLOP. It is a settled, independently verifiable unit of the promised workload at the promised tier.

DVC must incorporate:

- **scenario specificity:** drought, fuel shock, transformer loss, chip embargo, network partition, de-peg, and coordinated policy create different graphs;
- **substitution latency:** an alternative counts only if it enters before the service obligation expires; and
- **common-cause groups:** nominally separate facilities may share a firmware signer, transformer vendor, cable, gas market, identity layer, cooling basin, legal safe harbor, stablecoin, or market maker.

Sovereign Optionality asks how many independently substitutable paths exist. DVC asks how much usable service reaches the sink after a scenario binds. A system can score well on one and poorly on the other. Both should be reported, along with the active binding constraints and dependency groups.

Receipts are evidence for edge capacities; they are not capacity. Work Credit and capacity-voucher issuance must be bounded by a conservative stress result, computed jointly across workloads that share capacity,

$$Q_{W,t}^{\text{issue}} \leq \min_{s \in \mathcal{S}_{\text{constitutional}}} \{ \text{DVC}(W, t, s) \} h_{W,t}, \quad 0 < h_{W,t} < 1,$$

with a published prudential haircut $h_{W,t}$. Nameplate power, installed hardware, gross proof throughput, and benign averages are inadmissible substitutes.

6.2 The Pressure–Capacity Corridor

Let $A_{\text{full-stack}}(R)$ be complete-path availability under regime pressure and let $H(R)$ be the holder’s practical ability to self-custody, transact, prove, and exit. Potential monetary service is bounded by

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

The likely shape is an inverted U, not a crisis escalator. Under low pressure, institutional substitutes work and differential value is thin. Under moderate pressure, substitutes weaken while the open stack and holder remain operational. Under extreme pressure, power, networks, hardware, liquidity, custody, or exit can fail. More crisis does not imply more premium.

7 Native Buyer Quality: Demand Is Not an Anchor

A dollar of fee demand is not necessarily a dollar of monetary risk absorption. Users can acquire the asset immediately before paying a fee. Operators can sell receipts immediately. A burn removes supply but creates no buyer. Collateral can unwind during stress. A wrapper can create price exposure without self-custody or native use.

Link 8 therefore requires a material, persistent cohort that is:

- **self-custodied**, rather than represented only by custodial or synthetic claims;
- **loss-bearing**, funded without a margin call, daily reset, or fragile redemption promise;
- **long-horizon and persistent**, with holding periods that extend beyond transaction windows; and
- **regime-responsive and countercyclical**, accumulating through relevant drawdowns and pressure rather than only after recovery.

A balance that falls because it was spent on the service during stress is the mechanism working, not a holder leaving: the map classes every reserve outflow as forced liquidation, operating turnover, or exercise, and only the first counts against the cohort.

The Native Monetary Buyer Map must separate just-in-time fee acquisition; operator sell-through and inventory; burns and net issuance; self-custodied reserve accumulation; wrapper demand; leverage and rehypothecation; holding period; loss-bearing capacity; and countercyclical accumulation. Unknown holdings remain unknown, and unknown is a coverage limit rather than a quality grade: taking part in the measurement is not a condition of privacy. Pseudonymous addresses must not be assigned to a favorable category without evidence.

Buyer quality is a checklist, not a scalar. Persistence, horizon, liability match, loss-bearing capacity, and countercyclicity answer different questions. A persistent automatic buyer can still fail as an anchor if its rule absorbs less risk during a drawdown. The reserve asset of the incumbent system shows how: a benchmark rule orders duration it does not hold, futures deliver it, and leveraged funds financed in repo take the other side, so the bid that looks most stable is the one financed most fragiley. Where holders are agents, a cohort that is one operating policy at many addresses, or that no person can still revoke, is not an anchor either. “There is demand” is incomplete without “from whom, with what capital, under what mandate, at what leverage, against what liability, and with what ability to survive loss.”

8 Triad Coherence and Instrument Hierarchy

8.1 Triad Coherence Test

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

Architecture A—one native monetary asset.

Strongest direct value capture and shared liquidity; greatest reflexivity, governance scope, and cross-domain collateral risk.

Architecture B—neutral reserve plus service-specific credits.

Clean separation between reserve money and typed service claims; weaker new-token capture and a need for interoperable settlement and independent credit markets.

Architecture C—shared settlement plus modular domain collateral.

Common payment rail with service-specific risk containment; greater operational complexity, collateral fragmentation, and bridge-governance burden.

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

8.2 The hierarchy is constitutional

Evidence objects.

PIDL receipts and work receipts prove that an interaction or work unit occurred. Facility Energy and Facility Capacity Receipts provide evidence about infrastructure and edge capacity. They are copyable evidence, not money and not the capacity itself.

Capacity and service objects.

Work Credits are typed, prepaid claims on reserved future service, owed by a named provider or pool: capacity is committed, the credit is sold and the capacity reserved in one atomic step, and verified delivery extinguishes the credit and issues the receipt. Workload, location, hardware, SLA, delivery window, and remedies are written in the claim; cancellation, expiry, default, and refund follow its terms. They may be transferable, redeemable, or useful for procurement and hedging. They are not presumed savings.

Duration-bearing credit.

Project notes, capacity bonds, purchase agreements, and infrastructure tranches carry maturity, covenants, default states, and loss waterfalls. Proofs can audit them; they do not become monetary backing.

Derivatives and operating claims.

LP and staking positions expose the holder to fees, slashing, validation, or corridor operation. They are not monetary objects.

Conditional monetary candidate.

Only the separate bearer base asset is tested for monetary premium, and only after DVC, non-bypassability, buyer quality, agency, and the remaining chain conditions pass.

No par guarantee, redemption promise, emergency support, or project-credit loss may be transferred to the base asset. Money stays duration-neutral; credit stays labeled, underwritten, and loss-bearing. Proofs can reduce information asymmetry and expose covenant breaches. They cannot eliminate construction delay, commodity risk, seizure, obsolescence, default, or time.

9 Strongest Objections and What Is Not Proved

9.1 The objections that remain load-bearing

Why not just Bitcoin? Bitcoin has the superior monetary work function, demonstrated scarcity, and an established holder base. The response is not that AfterFiat improves Bitcoin. Portable attestation and verified compute are separate goods that Bitcoin does not attempt to supply. The open question is whether an asset associated with those goods earns monetary premium or merely prices as a service. Losing to Bitcoin can coexist with a successful protocol.

Utility does not imply money. Correct. Fees, burns, and collateral can make a valuable service asset. The monetary claim survives only in the holder-side state-contingent service, and only if the complete stack and holder remain usable under pressure.

The asset can be bypassed or copied. Correct unless the differential value exceeds verification overhead and the fee, and unless network effects such as liquidity, anonymity sets, collateral, operator relationships, workload standards, and receipt history make a lower-fee fork less attractive. Open specifications lower the cost of copying. This is an empirical moat, not a theorem.

Proofs do not prove truth. Correct. They verify bounded claims under stated assumptions: origin, custody, computation, inclusion, policy predicates, and settlement state. They do not establish honest inputs, appropriate models, or social meaning.

The closed sovereign stack wins. It may win on capability. The thesis depends on a constituency willing to pay for custody, privacy, portable verification, and exit rather than merely throughput. If users retain practical agency inside a competent integrated system, the open-stack premium narrows.

The asset is co-opted by wrappers. It may become a successful financial product while native monetary use stagnates. Wrappers can improve liquidity and legal holdability while centralizing custody and manufacturing holders without users. That outcome is measurement, not adoption.

Stablecoins are the incumbent bypass. A decade of revealed preference sits in this channel: in exactly the financial-repression states this thesis targets, dollar stablecoins—not bearer assets—were the observed instrument of capital flight, settling at scale on Tron and Ethereum with no native monetary asset in the loop. For the median settlement need the incumbent channel is cheaper, faster, and deployed. The differential survives only in the upper band—private settlement where a stablecoin is a custodial IOU, proofs where it is silent, verified compute where it has nothing to verify—and where the triad service is interchangeable with a stablecoin payment, non-bypassability should be presumed failed rather than defended.

The value-capture design carries classification risk. The strongest legal threat is classification, and it is examined per transaction, not per token. The Commission's March 2026 interpretation leaves *Howey* in force and separates whether an asset is a security from whether a transaction in it is subject to an investment contract; the Division's September 2026 staff FAQ, a staff view without legal force, treats a protocol-funded burn on a functional system with no central party as no promise of managerial effort. The open questions therefore sit with native issuance, secondary transfers, service prepayments, staking pools, and project finance, and they turn on governance concentration and discretionary return management rather than on value capture as such. Issuance and fee policy fixed ex ante beyond governance reach narrows the exposure without conferring clearance, and it is also the version closest to what a monetary constitution should be. The analysis is dated, and none of it is advice.

The joint failure is the one to watch. The closed-stack and wrapper-dominance threats are usually read separately; the failure that would arrive with every dashboard green is both at once—services from a competent closed stack, price from the wrappers, the base asset reduced to a reference price with no native loop while every dashboard reads green. The full edition names the diagnostic: persistently positive wrapper-native growth gap, no corridor increment in the convenience yield as regime pressure rises, and a falling homestead ratio against improving closed-stack

service metrics. The triple divergence is observable years before the individual red lines bind.

Who finances the reactor? Not the base asset. Long-lived infrastructure requires explicit project credit, underwriting, covenants, equity, and a holder willing to bear time. If the protocol turns Work Credits or the base asset into disguised construction bonds, it has re-entered the duration trap it was designed to avoid.

9.2 What this edition does not prove

It does not prove that:

- Privacy, Proofs, and Compute should share one asset;
- the native asset can sustain a material non-bypassable fee;
- fee demand or burns create durable reserve demand;
- Work Credits should be treated as savings instruments;
- an open stack remains physically available in the states where its differential value is greatest;
- the holder-side convenience yield is large, or larger than the premium already captured by gold or Bitcoin;
- a balance held in advance does better than Bitcoin converted on the day, for any workload or stress state;
- wrappers will not dominate price and custody;
- states will permit lawful privacy and practical exit at useful scale;
- useful-work markets can avoid hyperscaler, hardware, router, or governance concentration;
- decentralized markets can finance long-lived infrastructure without concentrated credit intermediaries;
- the median buyer will pay for cryptographic verification while courts and creditworthy indemnitors work; or
- agent demand, where it exists, is monetary, or answerable to a person.

Passing all nine links demonstrates only a defensible monetary candidate. It does not establish timing, valuation, inevitability, or a tactical view on bonds, currencies, commodities, or any geopolitical scenario.

It also carries a base-rate problem the edition states rather than buries: no *equity-like claim on a fee stream* is known to have crossed to monetary premium. Warehouse receipts, bills of lading, and standardized commodity claims acquired price and liquidity without ever acquiring monetary treatment. Commodity money itself—silver, salt, tobacco—did cross from use-value to money along the path Menger described, and that account is an argument *for* the possibility; the empty record is specifically for fee-stream claims. The thesis proposes that a regime now exists in which that crossing could first occur. What follows is the argument for the possibility, not a precedent.

Every prior restructuring of the information substrate arrived as an externality of adoption rather than as anyone's plan, which is why this thesis publishes instruments rather than forecasts.

10 Eighteen Red Lines

The monetary claim is retired when a red line persists without credible remediation. An isolated incident is not automatically fatal; the governing pattern is sustained breach plus failed recovery.

Thresholds must be declared before the event and may not be relaxed while the remediation clock runs.

1. **Verification Affordability Breaks.** VERIFYPRICE exceeds constitutional bounds for core workloads for three consecutive months without a credible remedy. Public verification has become a platform service.
2. **Refund Safety Breach.** An admissible corridor produces protocol-attributable losses—any party without a bounded-time no-loss exit—and is not automatically delisted and remediated.
3. **Verification Monoculture.** More than 70% of verification remains on one hardware profile, TEE vendor, or jurisdiction for three months.
4. **Telemetry Capture.** Receipt datasets become unavailable, unverifiable, or controlled by one party. The scoreboard has become theater.
5. **Fee Coverage Collapse.** Retained-fee coverage (burns excluded) remains below 10% of the security budget while more than 80% of workloads are speculative for twelve months.
6. **Value Capture Failure.** Triad use grows for twelve months while native fees, burns, collateral, and fee coverage do not. Service is bypassing the asset.
7. **Legal Incompatibility.** Lawful users in major jurisdictions cannot use privacy rails without unacceptable uncertainty for twelve months, and scoped disclosure patterns do not gain adoption.
8. **Governance Capture.** Governance can alter issuance, fee routing, or telemetry without hard constraints, timelocks, and appropriate supermajorities.
9. **Wrapper Dominance Becomes Monetary Substitution.** Custodial and synthetic exposure grows for two quarters while native fees, private settlement, collateral, and use stagnate or decline.
10. **Physical Infrastructure Opacity.** Facility evidence is unavailable or unverifiable for a material capacity share, or physical audit cost remains above threshold, for two quarters.
11. **AI Enclosure.** Top hyperscalers, closed TEEs, or one jurisdiction persistently control a threshold share of verified compute, converting the service market into a cloud IOU.
12. **Agency Failure.** Forced disclosure becomes systemic, or institutional use grows for twelve months while user-level agency and participation use cases do not.
13. **Energy Sovereignty Failure.** Network sovereign optionality remains below threshold for two quarters, or a threshold share of capacity sits under active curtailment or rationing against verification workloads.
14. **The Capturable Wedge Closes.** Either the service premium net of verification cost falls below the realized protocol take for two quarters, or real native fee turnover fails to grow over four quarters while physical verified throughput grows.
15. **Native Collateral–Capacity Spiral.** A pre-declared material drawdown persists; collateral falls below realized slashing exposure; operator exit exceeds threshold; stress DVC falls; released collateral expands liquid float; and core SLOs fail to recover in the remediation period.
16. **Convenience-Yield Null.** Over at least eight quarters spanning at least two regime transitions, the upper bound on the holder-side yield’s increment in the pre-registered pressure corridor, over the benign state, lies below the pre-registered minimum, and the test had the power to see that minimum. (Re-specified in v3.5 from a linear coefficient, which an inverted U could trip at a slope near zero while the mechanism was present.)
17. **Information-Sensitivity Reversion.** With a mature lender panel, cross-lender haircut dispersion

on verifiable positions, controlled for volatility and depth and read across independently maintained underwriting approaches, is not lower than on opaque comparables for four quarters, or widens relative to them in a stress episode while receipts stay available. A panel without five independent approaches reads unidentifiable, not green.

18. **Denomination Null.** Outsiders decline to reckon in the unit (weakens); protocol insiders' short-tenor obligations—prices fixed for 90 days or less—fall short of a 25% unit-denominated share at Phase III entry, or stay below it for four quarters after (kills). Sales on receipt are evidence, not a trigger; long-tenor obligations carry no fatal threshold. (Re-specified in v3.5: the line had read long-dated obligations, which move last, and tripped on a share that was merely flat.)

Concentration lines (3, 9, 11, 13, 15) carry maturity gates: before the Phase II entry gate they are published as watch indicators, since any launch network trips them. Every red line carries a numeric trigger and a persistence window in the full edition. A candidate nineteenth line, **Principal Concentration**, is published with placeholder thresholds, rated *weakens*, and not counted: it reads on native holdings that are one operating policy at many addresses.

Red Line 15 captures wrong-way risk:

drawdown → more units required → operator exit → DVC decline
→ collateral release → liquid float increase.

Pure-native collateral maximizes direct value capture and this reflexivity. Mixed collateral weakens native capture and improves prudential resilience. Native stake plus a separately capitalized insurance or resolution pool splits the functions but introduces its own funding and governance questions. The trade-off must be reported rather than wished away.

A failed remediation reclassifies the asset from store-of-value candidate to speculative or experimental. A separate service-good outcome is also possible: the stack works, capacity is consumed, and fees accrue, but no holder-side insurance signature appears. Then the honest classification is a verified-capacity service asset, not money.

11 Four-State Macro Model

AfterFiat is structural architecture, not a tactical Treasury short. The same system can pass among four states, and policy repair can extend incumbent arrangements for years.

State 1—Market repair succeeds.

Long yields are driven mainly by term premium and impaired duration absorption while long-dated inflation compensation remains contained. Buybacks, maturity changes, regulatory repair, or lower front-end rates restore function. Private substitutes regain credibility and the immediate repression premium for an open stack may narrow.

State 2—Physical inflation blocks repair.

Energy, refining, gas, power, shipping, cooling, transformers, or industrial bottlenecks lift long-dated inflation compensation. The financial system may need easing while the physical system argues for restraint. Layer-0 costs and neutral- asset demand can rise together.

State 3—Market repair fails and repression follows.

Liquidity support, issuance changes, or easing fail to restore a durable marginal buyer. Policy shifts from repair toward captive demand, preferential regulation, negative real returns, capital controls, or restricted exit.

State 4—Competent closed-stack stabilization.

The state successfully coordinates energy, industry, compute, identity, payments, and duration warehousing. Functionality and collateral stability improve while user custody and practical exit contract.

Quiet long-dated inflation compensation leaves State 1 available. A sustained rise associated with physical bottlenecks points toward State 2. Failed repair plus coercive balance-sheet rules indicates State 3. Demonstrated integrated delivery with shrinking user exit indicates State 4. These are observables and transition conditions, not dated forecasts. The binding quantity is the interest-rate risk to be absorbed in a period against the capacity of those who absorb it, not the debt ratio, which is why the model has states rather than a threshold: absorption that runs through leverage works without incident until it binds, and the long yield is not the series that shows it first. External scenario commentary contributes mechanisms—physical bottlenecks, marginal-buyer quality, and policy constraints—not facts the thesis requires to be true.

12 Minimum Viable Experimental Roadmap

The roadmap is ordered by dependency and designed to produce disconfirming evidence early.

12.1 1. Receipts and verification harness

Define a compact, chain-agnostic receipt schema binding claim hash, workload ID, circuit or model hash, proof commitment, service tier, timestamps, resource use, hardware profile, and provider signatures. Ship deterministic test vectors, reference verifiers on ordinary hardware, signed benchmark results, and archived raw data. Publish VERIFYPRICE distributions, failure rates, and adversarial-input behavior. If independent parties cannot reproduce the dashboard, stop.

12.2 2. One privacy corridor

Ship one non-custodial corridor before announcing many, and make it the one that exists: BTC↔XMR, whose adaptor-signature protocol is deployed and whose refund semantics are understood. BTC↔shielded-ZEC swaps do not yet exist and are a research item, not a launch target. Specify atomic completion, bounded timeouts, a no-loss exit for every party within bounded time, scoped disclosure (per-payment, epoch-rotated), and clear failure UX. Publish VERIFYSETTLE success, time to finality, refund outcomes classified by root cause, per-asset anonymity metrics, route concentration, and incident reports. Any protocol-attributable refund failure excludes the route.

12.3 3. One proof factory and canonical starter set

Stand up an open-admission proof factory for a small workload set: a generic proof, matrix multiplication in both a probabilistic (Freivalds) and a succinct SKU, model inference with honestly graded assurance, media provenance, and atomic settlement. Every acceptance criterion ships with an adversarial test suite, and the matrix-product bound is componentwise, because an aggregate bound admits INT8 execution and a single corrupted output. Support multiple proof backends so the experiment does not encode monoculture at launch. Measure queue depth, entry latency, verifier cost, prover concentration, and service tiers.

12.4 4. DVC and physical evidence

Create Facility Energy and Capacity Receipts for participating facilities. Build the energy-to-settlement graph, name dependency groups, measure substitution latency, publish benign and stressed DVC, and name each active binding constraint. Bind outstanding Work Credits to the conservative stressed result with a declared haircut, computed jointly across workloads that share capacity. Test transformer loss, hardware restriction, network partition, curtailment, common firmware failure, and liquidity disruption.

12.5 5. Developer surfaces and neutral routing

Only after the primitives work, expose simple verbs to declare a claim, request a proof, await a receipt, verify a receipt, and pay conditionally for proof or compute. Add neutral routing once there is enough activity to measure fairness. Publish time to first work for a new provider, house share, top-provider share, geography, ASN diversity, and rejection reasons.

12.6 6. Monetary and collateral experiments

Test the three Triad architectures rather than assuming one token. For any base asset, publish native fee share, burns, issuance, collateral, bypass prices, the capturable wedge, and the Native Monetary Buyer Map. Stress pure-native, mixed, and native-plus-insurance collateral against the same draw-down and DVC scenarios. Run the access experiment: equal budgets fixed before a disruption, a held native balance against Bitcoin bought on the day, provider-side conversion, and reserves plus service reservations, first as an analog on an existing network whose fees must be paid in its own asset. Keep Work Credits typed, reserved, and DVC-bounded; keep project finance in explicit credit instruments.

12.7 7. Useful-work pilots last

Only after receipts, corridors, proof markets, telemetry, and routing are stable should useful-work consensus be piloted. Candidate laboratories include matrix-multiplication work, verified-inference service, and a proof-receipt ledger. Their purpose is to test verification asymmetry, open entry, decentralization, and cheap checking at scale—not to announce a final monetary base. They also exist to study, not to assume away, the open problem at the centre of proof-of-useful-work: inputs derived from block headers are unpredictable and therefore useless to any buyer, while inputs supplied by clients are useful and therefore precomputable. A hybrid design with a stated,

tested collusion bound is the flagship research problem for any lab that builds this.

12.8 Decision rule

Advance by telemetry gates, not calendar promises. Before verification, settlement, reachability, DVC, value capture, and holder quality are public and healthy, the base asset is research exposure. A working service stack is infrastructure. Only repeated performance under real or simulated pressure, with a material native holder anchor and no red-line breach, permits store-of-value language. Even then, the premium remains unsized.

13 What Is Being Bought: Stack First, Coin Optional

The full edition names the institution it proposes—**Kardashev Labs**, a build-and-measure lab—and states what an allocator would be funding. The concessions above fix the answer. Bitcoin’s work function is the better money; utility-token precedents reached low-single-digit fee coverage; the fee base arrives, if at all, on the far side of a clearing event the thesis cannot time; the holder-side premium is unsized and instrumented for its own retirement.

The Investment Thesis

Fundable as public infrastructure: the verifiable stack, with a measurement institution attached—the VERIFYPRICE observatory and its adversarial harness, VERIFYREACH, VERIFYSETTLE on one deployed corridor, the PIDL receipt schema and SDKs, facility capacity receipts, and VERIFYFLOW wrapper telemetry. Each has prospective buyers who are not token holders; VERIFYFLOW against existing bitcoin wrappers could be sold before any base asset exists, and whether anything is sold is itself a kill condition for the lab.

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. Nothing in the lab’s budget depends on it, and nothing the lab ships is worthless if it fails.

Not on offer: a token launch, a protocol treasury, or early exposure to a base asset. The lab issues no monetary instrument.

The full edition supplies dated deliverables mapped to the minimum viable stack, a team shape, the structure of a planning envelope funded in deliverable-gated tranches (the figures live in a separate costed plan), an entity and licensing posture, the seven research problems the thesis names as open, and kill conditions for the lab itself. Every deliverable is a Phase I artifact, because under the thesis’s own model Phase III is not reachable on a 36-month horizon, and a plan that claimed otherwise would contradict the document it is attached to.

14 Conclusion: Equip, Do Not Manage

AfterFiat is strongest as a research and engineering agenda, not as a succession narrative. It asks whether private settlement, portable proof, and verified compute can be made into ordinary, stress-deliverable infrastructure, and whether holding a separate bearer base asset can preserve enough agency through that infrastructure to earn monetary premium.

The answer may be no. The triad may remain three useful service markets beside Bitcoin or another neutral reserve. A protocol can succeed while its base asset realizes as a service claim. A closed sovereign stack can deliver more capacity at lower cost. A holder-side premium can exist in mechanism and prove too small in magnitude. An open stack can be most needed in the same state in which its physical delivery fails.

Those are not footnotes. They are the load-bearing terms of the bet.

The base asset is only a conditional monetary candidate. Work Credits are typed service or capacity claims. Receipts are evidence. Project notes are explicit credit. LP and staking positions are derivatives or operating claims. No useful service becomes money by terminology, no price chart proves adoption, and no dashboard reading above its threshold may be explained away after the fact.

The mandate is therefore concrete: build receipts, rails, and verifiable machines; measure the full path from energy to holder utility; preserve self-custody and exit; keep credit separate; identify who bears monetary risk; and retire the claim when its published conditions fail. Privacy equips people to transact without default legibility. Proofs equip them to check claims without trusting a platform. Compute equips them to purchase machine work without a single permissioning chokepoint. Telemetry equips the public to see when any of those promises becomes another soft guarantee. Where people act through agents, equipping them means a key root they hold and a revocation they can actually use.

Gold condenses geology. Bitcoin condenses thermodynamics. The triad may condense verifiable agency in the band where institutions weaken but networks still work. The first two are established monetary facts. The third remains an experiment.