Bitcoin as Unique Neutral Settlement: A Seven-Property Elimination

Sean Hash
bitcoingametheory.com
sean@bitcoingametheory.com

April 2026

Working Paper. Cite as Hash (2026b).

JEL Codes: G11, G12, G15, E42, E51, L14

Keywords: Bitcoin, neutral settlement, asset comparison, digital scarcity, reserve assets, capital allocation, property elimination, game theory, network effects

Abstract

Which asset can serve as neutral settlement in a multipolar world? Hash (2026a) establishes that the payoff advantage of Exit over Stay is monotone increasing in adoption, and derives seven necessary properties: protocol security, neutrality, permissionless access, cheap finality, absolute scarcity, informational security, and adaptive resilience. This paper applies those properties. We conduct a systematic elimination across seven asset classes: fiat currencies, sovereign bonds, equities, real estate, gold, alternative blockchain tokens, and Bitcoin. Six fail at least one required property. Bitcoin satisfies all seven because protocol neutrality is independent of ownership concentration. Under proof-of-work with cryptographic custody, holding coins confers no consensus vote, no ability to alter supply, no ability to censor transactions, and no ability to freeze other holders' keys. Capture surfaces at ETF wrappers, exchanges, mining pools, and corporate treasuries are real, but each is unilaterally exitable without counterparty consent, and even a coalition controlling a majority of transferable supply could not prevent anyone holding keys from settling. No other asset separates ownership from protocol authority in this way. The analysis does not require fiat collapse. Bitcoin operates as settlement finality layer in parallel with fiat payment rails.

1. Introduction

The Exit Game (Hash, 2026a) proves that the payoff advantage of moving capital from capturable settlement systems to neutral settlement is monotonically increasing in adoption. But the framework is agnostic about *which* asset serves as the Exit destination. It derives seven necessary properties (P1–P7) and shows that the Exit advantage is monotone increasing. The identity of the neutral settlement asset is an empirical question, not a theoretical one.

This paper answers that question. We take P1–P7 as given and ask: which existing asset satisfies all seven simultaneously? The answer is reached by elimination, not by advocacy. Each asset class is tested against each property. Six classes fail. One survives.

A second contribution distinguishes this paper from Hash (2026a). The Exit Game treats capture as a binary, protocol-layer constraint. This paper relaxes that idealization. Every asset has capture surfaces at some layer, whether issuance, custody, venue, or physical possession. The relevant question is not whether capture exists but whether any feasible coalition can use that capture to alter settlement rules or prevent key-holders from transacting. We formalize this as Ownership-Protocol Independence (Proposition 3) and show that Bitcoin is the unique asset satisfying it. The argument is therefore comparative. Bitcoin is not uncapturable at every layer; its application-layer captures cannot translate into protocol-layer control under proof-of-work, because coin ownership confers no consensus authority.

The structure is deliberate. By separating the game-theoretic argument in Hash (2026a) from the asset identification in this paper, we make the logic auditable: a reader who rejects our claim about Bitcoin need only show that some alternative satisfies P1–P7, without engaging the game theory. A reader who accepts Bitcoin's properties need only verify the elimination table, without reconstructing the dominance proof.

2. The Property Framework

We summarize the seven properties derived in Hash (2026a), Section 3.3. Each property blocks a specific attack class that would violate neutral settlement, defined by Hash (2026a) as immunity from seizure, debasement, and political capture.

2.1 Formal Specifications

P1 (Protocol Security). Mining constitutes a Nash equilibrium; attack cost exceeds expected gain. Formally: $C_{\text{attack}} + P_{\text{collapse}} > \Delta_{\text{attack}}$, and $\sum \delta^t \cdot \mathbb{E}[\Pi_{\text{honest}}] > \Delta_{\text{attack}} - C_{\text{attack}}$ (Budish, 2018; Biais et al., 2019).

P2 (Neutrality). No issuer, no foundation, no governance mechanism capable of unilateral rule changes. $|CS_{\text{protocol}}| = 0$.
P3 (Permissionless Access). Any agent can initiate settlement without third-party permission. $\Pr(\text{censor access} \mid \text{scale}) \approx 0$.
P4 (Cheap Finality). Settlement of \$1 billion for fees that scale with transaction bytes, not value: under \$1 at prevailing mempool rates (1 sat/vB, April 2026) and under \$520 even at the highest sustained congestion on record (~1,200 sat/vB during the April 2024 halving; ~500 sat/vB during the May 2023 BRC-20 event), with mathematical finality in under 60 minutes.
P5 (Absolute Scarcity). Fixed supply with zero supply elasticity: $dS/dP = 0$.
P6 (Informational Security). Custody is mathematical (private key knowledge). Cost of seizure scales superlinearly with targets.
P7 (Adaptive Resilience). Protocol upgrades via consensus without introducing governance capture.

2.2 Necessity and Sufficiency

Proposition 1 (Necessity). *P1–P7 are individually necessary. Removing any single property enables the corresponding attack vector.*

Proved in Hash (2026a).

Proposition 2 (Sufficiency). *P1–P7 are jointly sufficient. Candidate additional properties either reduce to P1–P7 or introduce non-structural criteria that would create capture surface, violating P2.*

For example, "privacy" as a candidate property reduces to P6 (informational security).

2.3 Ownership-Protocol Independence

The seven properties in Hash (2026a) are stated as binary constraints at the protocol layer. At the application layer, every asset admits capture surfaces that P1–P7 do not directly address. We introduce a formal complement.

Definition 1 (Protocol Capture Fraction). For asset A, let $\kappa_{\text{protocol}}(A) \in [0, 1]$ denote the fraction of A's settlement rules, supply schedule, and transaction validity that any feasible coalition can unilaterally alter.

Definition 2 (Ownership Capture Fraction). Let $\kappa_{\text{owner}}(A) \in [0, 1]$ denote the maximum fraction of A's transferable supply any single coalition of state actors, custodians, issuers, or venue operators can hold, seize, freeze, or render non-transferable.

Proposition 3 (Ownership-Protocol Independence)**.** *Neutral settlement requires* $\kappa_{protocol}(A) = 0$. *For assets where consensus authority or governance rights are conferred by ownership,* $\kappa_{owner} \to 1$ *implies* $\kappa_{protocol} \to 1$*: equities via shares, proof-of-stake tokens via stake, governance tokens via votes, stablecoins via issuer control. For proof-of-work Bitcoin with cryptographic custody, the implication fails:* $\kappa_{protocol}(Bitcoin) = 0$ *for any value of* $\kappa_{owner}(Bitcoin) \in [0, 1]$.

The proposition is a necessary complement to P1–P7. An asset satisfying all seven protocol-layer properties still fails as neutral settlement if ownership concentration translates into consensus or issuance control. Conversely, an asset with high ownership concentration remains neutral if protocol rules are enforced independently of coin ownership. Under proof-of-work, consensus is determined by hashrate and node operation, neither of which is conferred by holding coins. A coalition holding 51% of transferable supply cannot alter the 21-million cap, cannot censor transactions it does not see, cannot seize coins held by other key-holders, and cannot change consensus rules without independent agreement from miners and nodes. The elimination in §3 applies P1–P7 first, then verifies $\kappa_{\text{protocol}}(A) = 0$ for the survivor.

3. Elimination

3.1 Method

For each asset class, we identify which properties are violated and the mechanism of violation. An asset fails the test if it violates *any* single property, since P1–P7 are individually necessary. Where ownership concentration translates directly into protocol control, we note it explicitly, since Proposition 3 makes this the critical failure mode for governance-token and equity-like assets.

3.2 Results

Table 1 summarizes the outcome of the elimination. Six of seven asset classes fail at least one necessary property; Bitcoin is the only candidate that violates none. The buckets can also be read as points along a capturability continuum: fiat and sovereign bonds at the most-capturable end (issuer-controlled, denomination-controlled); equities and real estate next (state-regulated corporate form, or immovable and title-registered); gold and alt-L1 tokens further out (portable-but-physical, or digital-but-governance-captured); and Bitcoin at the non-capturable extreme (informational custody, no issuer, no founder). The detailed analysis in §3.3 works through each row in turn.

Asset Class	Properties Violated	Mechanism
C1a: Fiat currencies	P2, P5	Central-bank issuer with discretionary supply schedule; $\lvert CS_{\text{protocol}} \rvert > 0$ by construction
C1b: CBDCs	P2, P3, P5, P6	Fiat failures plus programmable censorship and state surveillance
C2: Sovereign bonds	P5	Negative real yield under financial repression; denominated in debaseable currency
C3: Equities	P2	Regulatory capture via margin requirements, trading halts, and dilution; ownership above 50% of shares confers board control
C4: Real estate	P3, P4, P6	Registration-gated transfers, multi-week settlement, physically seizable, subject to property tax and eminent domain
C5: Gold	P4, P6	Settlement friction of 3–8% for insurance, transport, and assay; physically seizable
C6a: Alt-L1 tokens	P2	Governance capture via foundation treasuries, VC voting blocs, or proof-of-stake concentration
C6b: Stablecoins	P2, P5	Issuer freeze function; denominated in fiat (inherits P5 debasement)
C7: Bitcoin	None identified	None

3.3 Detailed Analysis

C1: Fiat Currencies and CBDCs. Central banks set monetary policy to serve domestic objectives (Hash, 2026a, Assumption 1). The supply schedule is discretionary: $dS/dP \neq 0$ by design. This violates P5 directly. The central bank is itself the issuer with unilateral authority to alter the supply schedule, which makes $\lvert CS_{\text{protocol, fiat}} \rvert > 0$ and violates P2. CBDCs inherit both failures and add programmable censorship (violating P3) and state surveillance (violating P6). No fiat currency has maintained purchasing power over a 50-year horizon against hard assets.

C2: Sovereign Bonds. Bonds are claims on future fiat. When the underlying currency is debased, bondholders bear the loss through negative real yields. Financial repression, holding interest rates below inflation, is the observed mechanism. This is not a market failure; it is policy working as intended. Real yields on 10-year US Treasuries were negative for 26 of 36 months from 2020–2023 (Federal Reserve Bank of St. Louis, 2024, FRED series DGS10 minus T10YIE).

C3: Equities. Productive assets generate real returns but are subject to regulatory capture. Governments can impose capital controls, windfall taxes, forced sales, trading halts, and beneficial ownership reporting requirements. The capture surface is large: $\lvert CS_{\text{protocol, equities}} \rvert \gg$

0, and ownership concentration above 50% of shares confers board control, collapsing the ownership and protocol layers into a single coalition.

C4: Real Estate. Property is the paradigmatic *seizable* asset. Eminent domain, property tax, and zoning regulation constitute permanent capture surfaces. Physical location is known (violating P6). Settlement requires title search, escrow, and legal process (violating P4). Cross-border real estate settlement is measured in weeks, not minutes.

C5: Gold. This is the hardest case and deserves extended treatment. Gold stands here for the broader class of energy-to-produce commodities (silver, platinum-group metals, industrial metals, and historically also shells, salt, and cattle); gold is the representative because it is the most durable and has the lowest and most predictable supply growth rate within the class. Real estate (C4) and art are also commodity-like: each requires energy or labor to produce and holds physical form. They differ from gold on portability. Real estate is immovable, which makes it maximally capturable by the political system that controls its jurisdiction; this is why it separates out as its own bucket with a harder P6 failure and an additional P3 failure via registration-gated transfers. Art is portable but heterogeneous and provenance-dependent, which raises verification cost above gold's and pushes P4 failures further. Everything argued below about gold applies *a fortiori* to other commodities, which add supply elasticity failures (P5) on top of gold's structural issues. Gold has served as neutral settlement for millennia. It fails on two properties: P4 on settlement friction and P6 on seizure. Physical gold settlement incurs 3–8% costs including insurance, transport, assay, and storage (BullionVault, 2025). The efficiency gap versus Bitcoin settlement is approximately four orders of magnitude at peak congestion and seven orders of magnitude at prevailing rates, reflecting the fundamental difference between physical and informational assets. Executive Order 6102 (1933) demonstrated that gold is seizable at national scale.

The verification game. Gold settlement embeds a verification game that has no Bitcoin analogue. Model the interaction between a seller (S) delivering an asset claimed to be gold and a buyer (B) who must decide whether to verify:

	S: Authentic	S: Forge
B: Trust	$(V - P,\ P)$	$(-V,\ V + P)$
B: Verify	$(V - P - C_V,\ P)$	$(-C_V,\ -C_F)$

where V is the asset value, P is the price, C_V is verification cost, and C_F is forgery cost. The equilibrium depends on the relationship between C_V and C_F.

For gold, C_V is **monotone increasing in forgery sophistication**: spray-painted lead is caught by visual inspection (~$0), but tungsten-core bars require destructive assay (2-5% of bar value) or ultrasound testing ($200-500/bar). Documented cases include tungsten-core bars in commercial LBMA inventories and spray-painted lead bars in retail markets. The US gold reserve at Fort Knox has not been independently audited since 1953, evidence that even sovereign-scale verification is prohibitively costly.

Proposition 4 (Verification Cost Asymmetry)**.** *Gold verification cost is strictly increasing in counterfeiter sophistication, while Bitcoin verification cost is asymptotically zero and invariant in attacker sophistication.*

Let $C_V^{\text{gold}}(s)$ be the cost of verifying gold against a counterfeiter of sophistication $s \in [0, 1]$, and let C_V^{btc} be the cost of verifying a Bitcoin UTXO. Three claims:

(i) $(C_V^{\text{gold}})'(s) > 0$: gold verification cost increases with forgery sophistication.

(ii) $C_V^{\text{btc}} = \varepsilon \approx 0$, independent of attacker sophistication, under Hash (2026a, Assumption 3).

(iii) Gold verification is non-persistent: each transfer resets the verification game, since the bar could have been swapped between inspections. Bitcoin verification is persistent: a confirmed UTXO remains valid until spent.

Proof. (i) follows from the physical structure of gold: detecting surface-level visual fakes costs less than detecting deep fakes via destructive assay. More sophisticated forgeries require more invasive and more expensive testing. (ii) follows from Hash (2026a, Assumption 3): verifying a digital signature is computationally trivial regardless of the attacker's resources (the attacker cannot produce a valid signature without the private key). (iii) follows from the difference between physical and informational assets: a physical bar's composition can change between inspections; a UTXO's validity is determined by the blockchain state, which is globally consistent and tamper-evident.

The implication for autonomous agents is categorical: an AI agent verifying gold must either trust a human intermediary, which introduces counterparty risk at $r = 0$, or deploy physical inspection infrastructure whose calibration itself depends on human trust chains. Bitcoin verification requires only a full node, software the agent controls entirely.

The gold-ETF objection. Gold ETFs settle digitally at low cost, arguably satisfying P4. But gold ETFs are claims on custodial gold, not gold itself. They introduce counterparty risk via the custodian and regulatory capture via the issuer. An ETF can be frozen, diluted, or restructured by its sponsor. If we accept ETF wrappers as satisfying P4, then by the same logic Bitcoin ETFs satisfy P4, and the underlying asset remains the object of analysis, not the wrapper. Gold's P4 and P6 failures are properties of gold-as-settlement, not gold-as-financial-product.

C6: Alternative Crypto Assets. This bucket covers non-Bitcoin blockchain assets, which divide into two subcategories.

C6a: Alternative L1 tokens. Every alternative blockchain protocol has at least one of: a foundation treasury that constitutes a governance capture surface, a venture capital allocation that creates concentrated voting power, or a proof-of-stake mechanism in which wealth concentration maps directly to protocol control. Ethereum's transition to proof-of-stake in September 2022 made this explicit: the largest stakers have proportional influence over block production and transaction ordering.

C6b: Stablecoins. Tokenized fiat and Treasury claims settle digitally at low cost and denominate in familiar units, arguably satisfying P4. But stablecoins fail P2 and P5 by construction. The issuer maintains a freeze function: Tether has frozen hundreds of addresses at the direction of law enforcement agencies, totaling over $1 billion as of 2025. This means $|CS_{\text{protocol, stablecoin}}| > 0$: a single entity can unilaterally alter settlement outcomes for any address. Stablecoins also fail P5: they are denominated in fiat currency, so the underlying asset debases when the reference currency debases, and the holder bears debasement risk with none of the yield. The capture is not hypothetical: proposed US stablecoin legislation (2025–2026) would cap yield payments to holders below market rates, a provision secured through banking industry lobbying to protect deposit bases. This is the P2 failure mode in its clearest form: a political coalition altering the asset's rules to serve incumbent interests, precisely the capture that neutral settlement requires immunity from.

C7: Bitcoin. Bitcoin satisfies all seven properties at the protocol layer. P1 is supported by the mining Nash equilibrium (Biais et al., 2019). P2 holds because there is no issuer, no foundation, and $|CS_{\text{protocol}}| = 0$. P3 is demonstrated by the network's continued operation across every jurisdiction that has attempted to ban it. P4 is empirically verified: $1B settles for under $500 even at peak congestion. P5 is algorithmic: a 21 million cap enforced by node consensus. P6 follows from cryptographic custody. P7 is demonstrated by the soft fork upgrade path without governance capture.

Application-layer capture surfaces exist and should be named directly. Spot ETF wrappers hold approximately 6% of transferable supply; public-company treasuries hold approximately 4%; exchange reserves hold approximately 8%; miner balances are negligible at the supply scale. Each captured fraction is unilaterally exitable: holders can withdraw ETF shares into spot Bitcoin and then into self-custody, miners can switch pools or solo-mine, and node operators can reject invalid chains. More importantly, under Proposition 3 these fractions are not load-bearing. Even if ETFs, treasuries, exchanges, and sovereign reserves in aggregate exceeded 51% of transferable supply, protocol neutrality would still hold. Holding coins does not confer the right to alter the supply schedule, censor transactions, seize other holders' keys, or change consensus rules. The 21-million cap is enforced by node operators running the reference implementation, not by coin holders voting their balances. This is the structural distinction between Bitcoin and every asset it competes with as reserve settlement.

3.4 Coexistence

The elimination does not require fiat collapse. Bitcoin operates as superior collateral and settlement finality in parallel with fiat payment rails. The cascade pressure (Hash, 2026a, Proposition 2) applies to the reserve settlement function. Fiat retains payment and unit-of-account roles, potentially indefinitely. This is not a revolutionary claim; it is a structural one.

4. Empirical Support

Theory without data is speculation. The following observations are consistent with the framework.

4.1 Institutional Adoption

Spot Bitcoin ETFs accumulated \$61.5 billion in net inflows since January 2024, reaching approximately \$170 billion in total AUM. BlackRock's IBIT crossed \$97 billion in 435 days, faster than any ETF in history (The Block, 2025). This is consistent with the cascade dynamics predicted by the Exit Game: institutional actors crossing their adoption thresholds p_i^* in clusters.

4.2 Sovereign Behavior

El Salvador holds 7,529 BTC with daily purchases and geothermal mining (Bitcoin Treasuries, 2025). Singapore issued 13 digital payment token licenses in 2024 (Monetary Authority of Singapore, 2024). Argentina received \$91.1 billion in cryptocurrency value amid 143% inflation (Chainalysis, 2024). These observations are consistent with sovereign defection from the Stay coalition (Hash, 2026a, Theorem 2).

4.3 Settlement Efficiency

Bitcoin settles \$1 billion in value for fees under \$1 at prevailing mempool rates (1 sat/vB, April 2026) and under \$520 at the highest sustained congestion on record (Mempool.space, 2024–2026), with mathematical finality in under 60 minutes. Gold settlement incurs costs across two dimensions that Bitcoin avoids.

One-time physical transport of LBMA-grade bullion at wholesale costs 0.5–1% of value for insured armored delivery (BullionVault, 2025), approximately \$5M–\$10M on \$1B, before assay at \$200–500 per bar across roughly 1,240 bars at 400 oz each, and before vault setup. Retail and cross-border settlement incurs 3–8%.

Ongoing ETF custody, the only format in which gold settles digitally, runs 17–40 basis points annually (SGOL 0.17%, IAU 0.25%, GLD 0.40%), or roughly \$2.5M per year on \$1B at the institutional median. Bitcoin self-custody has no recurring protocol cost.

At wholesale peak congestion, the one-time settlement ratio is approximately four orders of magnitude (\$5M / \$520 $\approx$ 10,000×); at prevailing rates it exceeds seven orders of magnitude. The recurring custody comparison has no Bitcoin analogue: ETF custody is a permanent fee stream that also introduces counterparty surface (P2 failure) and jurisdictional seizure risk (P6 failure). This is the empirical basis for the P4 elimination of gold.

4.4 Network Resilience

China's comprehensive mining ban in 2021 resulted in total hashrate recovery to all-time highs within six months (Cambridge Centre for Alternative Finance, 2022). Bitcoin's hashrate reached approximately 700 EH/s in 2025, with no PoW competitor exceeding 1% of this security budget (Cambridge Centre for Alternative Finance, 2025).

4.5 Historical Uniqueness

Bitcoin's position derives from an unreplicable historical sequence. First, pure origin: no ICO, no pre-mine, no venture capital, no foundation treasury (Nakamoto, 2008). Second, founderless development: the creator departed, with an estimated 1.1 million BTC unspent since 2009 (Lerner, 2019). Third, survival of six categories of existential crisis without protocol failure. Fourth, PoW hashrate dominance following Ethereum's migration to proof-of-stake in 2022.

No competitor can replicate this sequence. As Huberman, Leshno, and Moallemi (2021) observe, Bitcoin is a "monopoly without a monopolist," and the monopoly derives from the *absence* of a monopolist, not despite it. A new protocol with a known founder, a foundation treasury, and venture backing fails P2 by construction, regardless of its technical merits.

5. Discussion

5.1 The Gold Question

Gold is the strongest competitor and deserves extended treatment. Gold has served as neutral settlement for millennia. Its failure on P4 and P6 is not a deficiency of gold but a consequence of its physicality. The information revolution created a new possibility: settlement that is mathematical rather than physical, where custody requires knowledge rather than location.

Krugman (2018) argues that Bitcoin lacks fundamental value because it has no tether to economic activity. This is reasonable if you weight intrinsic yield over settlement properties. The present analysis weights structure, and on structure, gold fails two properties that Bitcoin satisfies.

Taleb (2021) makes a different case: that Bitcoin is fragile because it depends on continuous mining, electricity, and internet infrastructure. This is the steelman for P7 failure. Our response: seventeen years of operation through six categories of existential crisis (exchange failures, regulatory bans, mining exodus, protocol disputes, market crashes, and sustained negative press) constitute empirical evidence for P7, though not proof. The fee-transition risk (Carlsten et al., 2016) remains the strongest version of Taleb's concern.

5.2 Attack Survival

The elimination holds only if Bitcoin actually satisfies P1–P7. The most serious challenge is sovereign attacks: can state actors ban, seize, or suppress Bitcoin?

The steelman: a coordinated coalition of G7 plus China simultaneously criminalizing possession, enforcing through ISP-level filtering, and imposing secondary sanctions. The defense rests on three structural features. Informational security (P6) transforms seizure into a key-management problem. Permissionless access (P3) redirects activity rather than eliminating it, as China's 2021 ban empirically demonstrated. And the coalition itself is unstable: the first defector captures fleeing capital, and coordination problems are precisely what this paper series is about.

We regard F1, a permanent single global coordinator, as the most plausible falsification path. Not because it is likely, but because it is the only attack that does not face a game-theoretic self-undermining dynamic.

5.3 The Volatility Objection

Bitcoin's annualized volatility has exceeded 70% in multiple calendar years, with drawdowns exceeding 50% from all-time highs occurring in 2014, 2018, 2022, and again in early 2026 at approximately 52% decline from cycle highs. This is the most common objection from traditional finance practitioners and deserves direct treatment.

Volatility is not among P1–P7. This is deliberate: volatility is a market phenomenon, not a structural property of the protocol. The Exit Game model (Hash, 2026a, condition M2) posits that $\sigma'_B(p) < 0$: deeper markets reduce volatility. The long-term trend is consistent with this: Bitcoin's 90-day realized volatility has declined from roughly 150% in 2011 to roughly 45% in 2024. However, the empirical evidence is mixed. Sharp drawdowns persist at current adoption levels. The power law channel, under which price follows a power law of time since genesis, describes the central tendency but does not constrain short-term deviations.

We therefore distinguish two claims. The *structural* claim: volatility affects adoption timing through the $\lambda_i \cdot \sigma_B(p)$ term but does not affect equilibrium structure. More risk-averse actors have higher thresholds p_i^* and wait longer. This holds regardless of σ_B's sign. The *empirical* claim: volatility will decrease with adoption as markets deepen. This is plausible but unproven at current adoption levels, and the 2026 drawdown is a direct counterexample to the naive version. M2 should be understood as a long-run tendency, not a monotone empirical fact at every time horizon.

5.4 Settlement vs. Acceptance

The elimination framework rests on a distinction that deserves explicit treatment: *settlement* is not *acceptance*.

Settlement in Bitcoin is a cryptographic fact. A valid signature transfers value from address A to address B. The protocol does not ask who owns the addresses, whether the transaction serves a lawful purpose, or whether the receiving party will be able to spend the output. At the protocol layer, P3 (permissionless access) holds unconditionally.

Acceptance is a social fact. If address B is associated with a sanctioned entity, a designation maintained by OFAC, the EU, or other authorities, then exchanges and custodians will refuse to credit incoming funds from B or any address that received value from B. The coins are settled but economically impaired: their acceptance by regulated counterparties depends on off-chain arbitration that the protocol cannot control.

This creates a race condition. Mixing services, coinjoins, and cross-chain bridges attempt to break the provenance chain. Compliance tools (Chainalysis, Elliptic) attempt to trace it. The effectiveness of both evolves continuously. Partial laundering creates a spectrum of "cleanliness" rather than a binary clean/tainted classification, and gray-market exchanges accept coins that regulated exchanges reject.

For our analysis, the implication is: P3 guarantees settlement at the protocol layer, but the *economic utility* of that settlement depends on the post-settlement acceptance environment. The elimination holds for settlement functionality, since Bitcoin is the unique asset where the *protocol* imposes no barriers. Whether the *ecosystem* around the protocol adds barriers is a compliance question, not a settlement one. This distinction is sharpened in Hash (2026c), where autonomous agents face the compliance problem in its purest form.

The Acceptance Game. The settlement-acceptance distinction produces a natural enforcement equilibrium without requiring protocol-level censorship. Consider an actor receiving coins of uncertain provenance. The expected payoff of blind acceptance is:

$\mathbb{E}[\text{Accept}] = V - p \cdot d(t) \cdot c$

where V is the value of the coins, p is the probability of taint, $d(t)$ is the probability of detection at time t, and c is the cost of punishment (legal penalties, reputational damage, asset seizure). Acceptance is rational when $V > p \cdot d(t) \cdot c$.

Three features distinguish this game from the cash analogue. First, $d(t)$ is monotonically increasing: Bitcoin's permanent ledger means chain analysis capabilities improve over time and apply retroactively to past transactions. Coins accepted today may be flagged in 2030. Cash has no such property: once spent, provenance is lost. Second, regulated entities such as exchanges and custodians face c values that dominate V for any realistic taint probability, with license revocation, criminal liability, and institutional collapse on the downside, and these entities control the fiat on/off ramps. Third, the resulting "provenance discount" on tainted coins reduces the economic payoff of the underlying crime: you can steal Bitcoin, but you cannot easily convert it at par value.

The equilibrium stratifies by actor type:

Actor	d	c	Filtering
Regulated exchange	High	Very high	Mandatory
Individual (KYC'd)	Medium	High	Incentivized
AI agent	Medium	Low	Rational
Gray market	Low	Low	Minimal

Notably, the actors most capable of filtering (regulated exchanges) are also the actors most severely punished for failure, creating robust enforcement at the fiat conversion bottleneck. For autonomous agents, filtering converges to near-perfection not from moral obligation but from the trivial cost of on-chain provenance analysis relative to any nonzero expected punishment.

The policy implication: protocol-level censorship, which would violate P2 and P3, is unnecessary for enforcement, because the application layer produces its own enforcement equilibrium through rational acceptance decisions. Bitcoin's settlement-layer neutrality is compatible with discretionary filtering at the acceptance layer, and indeed is strengthened by it.

5.5 The Eighth-Category Objection

The elimination tests seven asset classes. A reader may ask: what about art, physical commodities, intellectual property, collectibles, carbon credits, private equity, tokenized real-world assets, or a future AGI-launched settlement protocol? The answer is reduction: each candidate collapses structurally into C1–C6.

Art, collectibles, wine, and watches are C4-family: physical, seizable, assay-dependent, illiquid. Physical commodities such as oil, wheat, and industrial metals are C5-family with one structural worsening: supply is elastic ($dS/dP > 0$) because higher prices draw out more production, so they additionally fail P5. Patents, trademarks, copyrights, carbon credits, and regulatory allowances are C1-family: the state is both the issuer and the enforcer, making $|CS_{\text{protocol}}| > 0$ by construction, with no exit path since protocol authority and ownership are fused. Private equity and venture stakes are C3-family. Crypto-collateralized stablecoins and governance-token protocols are C6: token-weighted votes over collateral, parameters, or issuance are foundation-equivalents. Tokenized real-world assets inherit the underlying's failures plus an issuer-wrapper P2 failure.

The sharpest form of the objection is an AGI-launched successor protocol. The game-theoretic answer is that a rational coordinator, whether human or machine, selects the protocol that already satisfies P1–P7. Proof-of-work Lindy, absence of a known founder, fair emission history, small block size preserving cheap full-node verification, and seventeen years of adversarial survival are path-dependent credentials that a freshly launched protocol cannot replicate. A new protocol with identifiable founders or a pre-mine fails P2 at genesis; one without them inherits Bitcoin's failure modes with none of its history. AGI actors face the same payoff structure as human actors in Hash (2026a): use what works, and fund its security through otherwise-stranded energy. Launching a competing protocol reintroduces the coordination failure the protocol was adopted to escape. Hash (2026c) develops the full game theory of Bitcoin neutral settlement under autonomous-agent adoption.

Even granting an unreduced edge case, the reserve-settlement function requires stock sufficient to absorb sovereign-scale flows without price dislocation that would destroy the reserve property itself. Art, collectibles, and intellectual-property markets are orders of magnitude smaller in transferable stock than the aggregate reserve base they would need to substitute for. The elimination is therefore robust to the eighth-category objection: the absent category either reduces to an already-failed bucket, or is irrelevant at the scale the question concerns.

5.6 Limitations

The elimination is as strong as the property framework. If P1–P7 are incomplete, meaning there exists an eighth necessary property that Bitcoin violates, the elimination fails. The sufficiency argument (Proposition 2) addresses this, but sufficiency proofs are inherently harder than necessity proofs because they require exhaustive enumeration of alternatives.

On minimality and completeness. We claim that P1–P7 are individually necessary and jointly sufficient. We do not claim they are *minimal*, meaning that no proper subset of P1–P7 suffices. Formal minimality would require showing that for each property P_k, there exists an asset satisfying all P_j for $j \neq k$ that fails as neutral settlement. The counterexamples are suggestive (Ethereum for P1, Solana for P2, CBDCs for P3, gold for P4, commodity bonds for P5, Swiss banking for P6, Litecoin for P7), but a rigorous minimality proof in the tradition of axiomatic characterizations in social choice (Arrow, 1951) and mechanism design (Gibbard, 1973) would require 7–10 additional pages and is not attempted here. The relevant question for the elimination is necessity plus sufficiency: can an asset fail any single property and still serve as neutral settlement? The answer is no, and that is what the framework proves.

Seventeen years of data is brief for a monetary claim. Gold has millennia. The response, that the relevant properties are structural and not historical, is logically sound but empirically untested at the timescales that matter. Carlsten et al.'s (2016) analysis of the fee-only security budget is particularly relevant: the transition from block subsidies to fee-only security is the largest untested structural assumption, and it will not be resolved by theory alone.

6. Conclusion

Global wealth is held in seven asset buckets: fiat currencies, sovereign bonds, equities, real estate, gold, alternative crypto assets (alt-L1 tokens and stablecoins), and Bitcoin. Neutral settlement requires seven properties: protocol security, neutrality, permissionless access, cheap finality, absolute scarcity, informational security, and adaptive resilience. The two sevens are independent. One partitions where capital sits, the other specifies what reserve settlement requires.

Six of the seven buckets fail at least one of the seven properties, and the failures are constitutive rather than contingent: a fiat currency without a discretionary issuer is not a fiat currency; gold without physicality is not gold; a smart-contract platform without a founder or treasury did not historically happen. Bitcoin fails none at the protocol layer. At the application layer, it is the unique asset whose captured fractions cannot translate into protocol control.

The claim is not that Bitcoin dominates on any single dimension. Gold is older, equities yield more, fiat is more liquid for daily payments. Nor is the claim that Bitcoin is uncapturable. ETF wrappers, exchange custody, corporate treasuries, and mining pool concentration are real capture surfaces, and each continues to grow. Ownership concentration may one day exceed 51% across institutional allocators. The claim is that none of this alters protocol neutrality, because under proof-of-work with cryptographic custody, holding coins confers no consensus vote and no ability to prevent other key-holders from settling. This is the structural difference Proposition 3 formalizes, and it is the difference that Hash (2026a) takes as given without proving. Hash (2026a) establishes why capital flows toward neutral settlement; this paper identifies which asset satisfies neutrality under partial capture; Hash (2026c) and Hash (2026d) treat the autonomous-agent and multipolar-sovereign regimes in which that flow accelerates.

References

Biais, B., Bisiere, C., Bouvard, M., and Casamatta, C. (2019). The blockchain folk theorem. *Review of Financial Studies*, 32(5), 1662–1715. https://doi.org/10.1093/rfs/hhy095

Bitcoin Treasuries (2025). Sovereign Bitcoin holdings tracker. bitcointreasuries.net.

Budish, E. (2018). The economic limits of Bitcoin and the blockchain. *NBER Working Paper* No. 24717. https://doi.org/10.3386/w24717

BullionVault (2025). Gold settlement cost analysis.

Cambridge Centre for Alternative Finance (2022). Cambridge Bitcoin Electricity Consumption Index: China mining ban impact.

Cambridge Centre for Alternative Finance (2025). Cambridge Bitcoin Electricity Consumption Index: Global hashrate data.

Carlsten, M., Kalodner, H., Weinberg, S. M., and Narayanan, A. (2016). On the instability of Bitcoin without the block reward. *Proceedings of the 2016 ACM SIGSAC*, 154–167. https://doi.org/10.1145/2976749.2978408

Chainalysis (2024). The 2024 Geography of Cryptocurrency Report.

Federal Reserve Bank of St. Louis (2024). FRED Economic Data: 10-Year Treasury minus inflation expectations (DGS10, T10YIE series).

Hash (2026a). Bitcoin exit dominance in monetary coordination games. SSRN Working Paper 6299081. https://ssrn.com/abstract=6299081.

Hash (2026c). Settlement at zero trust: Bitcoin and autonomous economic agents. Working paper, bitcoingametheory.com.

Hash (2026d). Monetary predator-prey dynamics: Enforcement gridlock and neutral settlement survival. Working paper, bitcoingametheory.com.

Huberman, G., Leshno, J. D., and Moallemi, C. (2021). Monopoly without a monopolist. *Review of Economic Studies*, 88(6), 3011–3040. https://doi.org/10.1093/restud/rdab014

Krugman, P. (2018). Transaction costs and tethers: Why I'm a crypto skeptic. *The New York Times*, July 31.

Lerner, S. (2019). The return of the deniers and the revenge of Patoshi.

Mempool.space (2024). Bitcoin transaction fee and settlement data.

Monetary Authority of Singapore (2024). Digital Payment Token Licensing Statistics.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.

Taleb, N. N. (2021). Bitcoin, currencies, and fragility. *Quantitative Finance*, 21(8), 1249–1255. https://doi.org/10.1080/14697688.2021.1952702

The Block (2025). Bitcoin ETF AUM tracker.

Notation

Symbol	Definition
P1–P7	Seven necessary properties: Protocol Security, Neutrality, Permissionless Access, Cheap Finality, Absolute Scarcity, Informational Security, Adaptive Resilience
$\|CS_{\text{protocol}}\|$	Protocol-layer capture surface: parameters an entity can unilaterally alter without network consensus
$\kappa_{\text{protocol}}(A)$	Fraction of A's settlement rules, supply schedule, or transaction validity any feasible coalition can alter
$\kappa_{\text{owner}}(A)$	Maximum fraction of A's transferable supply any coalition can hold, seize, freeze, or immobilize
dS/dP	Supply elasticity ($dS/dP = 0$ under absolute scarcity)
C_{attack}	Cost of attacking the protocol
P_{collapse}	Probability of network collapse conditional on attack
Δ_{attack}	Attacker's gain from a successful attack
Π_{honest}	Per-period profit from honest mining
δ	Discount factor
V, P	Asset value, transaction price (Verification Game, §3.3 C5)
C_V, C_F	Verification cost, forgery cost (Verification Game)
$s \in [0,1]$	Counterfeiter sophistication
ε	Near-zero computational verification cost for Bitcoin UTXOs
$d(t)$	Probability of taint detection at time t (Acceptance Game, §5.4)
p, c	Taint probability, punishment cost (Acceptance Game)
λ_i	Actor i's risk aversion
$\sigma_B(p)$	Volatility of neutral settlement asset at adoption level p
$\sigma'_B(p) < 0$	Volatility decreases as adoption deepens (Condition M2)
p_i^*	Actor i's adoption threshold