Orange Anchor

Bitcoin-Anchored Cost-Backed Claims and Verifiable Attribution

Abstract

Digital claims face a structural cost problem: once infrastructure exists, the marginal cost of producing additional claims approaches zero. This makes fake participation economically rational at scale and degrades the integrity of attention, governance, reputation, and identity systems. The cost is hidden but ubiquitous: businesses pay it as defensive spend, fraud loss, and lost revenue; users pay it as friction, throttling, intrusive verification steps, and service quality that quietly erodes year after year. Until now, the edge of the internet has had no practical way to enforce scarcity without relying on central intermediaries.

Online platforms, services, and communities lose significant resources every year to coordinated fake accounts that inflate engagement metrics, drain expensive services at industrial scale, and distort marketplaces, reputation systems, and public discourse.

We propose Bitcoin-Anchored Cost-Backed Claims. The approach produces digital claims that carry real production cost as a permanent property. Creating one costs real resources - time, energy, and memory - while verification stays fast, and independent of trusted third parties, all anchored to Bitcoin.

The construction restores positive marginal production cost to digital claims.

The first concrete construction is Orange Anchor, which runs on ordinary phones. A user's device performs sustained work during a real-time interval bracketed by Bitcoin block hashes. The resulting commitment carries real production cost as a permanent, verifiable property that travels with the claim across every system it is used in.

Honest users pay the cost once on hardware they already own and receive a reusable credential they fully control, with no ongoing production cost after the initial commitment. Producing fakes at scale can become economically unprofitable because every fake claim costs the attacker the full price, with no economies of scale.

Verifiers and businesses benefit from a higher-quality group of participants. Accepting these commitments as backing works like collateral or a security deposit. If someone behaves badly, the commitment can be publicly marked on the Bitcoin-anchored attribution record, lowering its future value as backing. This gradually reduces the population of fake actors, lowering defensive costs and improving service quality over time.

This changes the economics of fake participation from near-zero marginal cost to a positive linear cost structure. The pattern is reusable wherever claims need to be grounded in real resource expenditure.

The approach composes with the existing identity and trust infrastructure rather than replacing it. Decentralised identifiers, verifiable credentials, reputation systems, and governance mechanisms can each gain cost-backed assurance. The broader effect is structural: as more systems accept commitments as backing, the cost of operating fake actors at scale rises. Orange Anchor is the first concrete construction of the pattern.

1. The Problem

In physical systems, each additional claim requires its own resource expenditure. In digital systems, the infrastructure is built once and supports an unbounded number of subsequent claims, so the marginal cost per claim approaches zero through amortisation. This is a structural property of any system where replication dominates production.

A Bitcoin-anchored architecture for constructing self-issued digital objects with persistent verifiable production cost is the response this paper develops.

When marginal cost approaches zero, synthetic production at scale becomes economically rational for any actor who benefits from over-representation. The accumulated synthetic volume diminishes the signal integrity of legitimate participation across these systems.

Large-scale production of claims by a single controlling entity is known as the Sybil attack [1]. Where marginal production cost is near zero, Sybil attacks are economically rational at whatever scale the attacker's infrastructure permits.

Existing defences address volume rather than the underlying condition. Detection identifies claims as fake after production and requires a trusted judge. Gatekeeping restricts production to authorised parties and requires a trusted issuer. Both depend on trusted parties and leave the structural condition - near-zero marginal cost - unchanged.

Prior proposals have attached positive per-claim cost. Hashcash [2] used per-message proof-of-work to raise the cost of mass email. Proof-of-personhood schemes [3], including systems such as Worldcoin [4], BrightID, and Idena, combine cost or scarcity with identity claims. Each addresses one or two of the conditions - positive per-claim cost, non-transferability, permissionless self-issuance, anchoring to a public sequence, cost preservation across systems - without combining all five in a single self-issued, permissionless form. To our knowledge, no prior construction surveyed here combines time-binding to a public sequence, memory-hardness, and per-instance physical witnessing in a self-issued, permissionless form anchored to Bitcoin's specific block sequence. This paper proposes such a construction.

2. Cost-Backed Participation

Cost-backed participation already exists in the physical world. A regulated operator, a manufacturer, or a Bitcoin miner commits substantial capital and ongoing energy to an activity whose value depends on continued legitimate operation. The commitment's value is contingent on standing, not on the standalone market value of the asset. The capital is not consumed on entry; it remains at risk. If the operation loses legitimate standing, the hardware and energy are not transferred to anyone - they simply lose the value they previously held. The cost is real, sunk, and not recoverable by transferring the asset.

Three properties matter. First, the cost is real - capital and energy committed in ways that cannot be cheaply simulated. Second, the value of the commitment is contingent on continued legitimate participation. Third, the commitment is operationally non-transferable; the asset may sometimes be sold, but the legitimacy that gave it value cannot. Together these mean bad behaviour does not require confiscation to be costly: the cost is already sunk, and loss of standing destroys the value of what was committed.

The Orange Anchor construction restores these properties digitally. The collateral is a cryptographic commitment (termed the *envelope* in the *Orange Anchor Lexicon v2.7*) produced by a specific device during a defined interval, anchored to Bitcoin. Producing it requires sustained resource expenditure - time, computation, memory, and energy - over an interval bracketed by Bitcoin block confirmations; the interval itself is termed a *burn* in the Lexicon, and is available at two duration tiers (lightweight, 6 blocks; heavyweight, 144 blocks). The specific duration is a calibration parameter developed in the *Bitcoin-Anchored Collateral Commitments* construction paper (BACC v1.9). The interval cannot be compressed by capital, and the commitment cannot be reproduced by another device without that device performing equivalent work in its own interval. The construction therefore produces *persistent attributable scarcity*: cost that is real, sunk, and verifiable at production, and that travels with the commitment across every system and every interaction over its useful life.

Observations about a commitment holder can be recorded against the commitment in a public, Bitcoin-anchored record. Anyone may publish; anyone may read and interpret according to their own policy. No authority pronounces a commitment good or bad. The primary economic defence is the production cost itself; the public record is a reinforcing secondary layer.

3. The Substrate

Bitcoin [5] maintains a public, cryptographically linked sequence of blocks secured by the cumulative work of its network. Five properties matter:

- **Global.** Every participant sees the same chain.
- **Neutral.** The protocol treats all valid transactions identically.
- **Permissionless.** No authority controls reading or writing.
- **Permanent.** Entries at sufficient depth cannot be rewritten without expending energy proportional to all work since their confirmation.
- **Ordered.** Before-and-after relationships are globally unambiguous.

These properties define a *commitment surface*: a substrate where commitments can be recorded and later verified without trusting anyone to maintain or interpret them. A digital implementation of cost-backed participation needs exactly this - a record that is real, public, and not subject to unilateral revision. Bitcoin provides it.
BTC, the token, is the access mechanism: writing costs fees, and the fee market allocates scarce write capacity without a gatekeeper. The token's monetary properties are separate from the substrate's properties; this paper is concerned with the substrate. The architecture uses Bitcoin only for what it already provides - no protocol changes, no new token, no novel script. Commitments are recorded via standard transactions using established mechanisms for embedding application data. The footprint is minimal.

The architecture's use of Bitcoin is distinct from aggregation-style timestamping. Existing systems such as OpenTimestamps prove that data existed at or before a particular block, but do not surface the per-block hashes consumed during an interval as live computation inputs. The construction proposed here requires those specific hashes as inputs to the work performed during the interval, not merely their existence as timestamps. The first anchor establishes the logical start of the interval; the final anchor confirms the last block hash actually consumed during the work.

4. The Architecture

The architecture has two phases with asymmetric cost. In the **commitment phase**, real resources are expended on a specific device during a defined interval, and the resulting transcript is anchored to Bitcoin. In the **verification phase**, that transcript is checked against Bitcoin's public record by any party with the chain, without contacting an issuer, platform, or authority. Bitcoin is the unchangeable record between the two. (The *Orange Anchor Lexicon v2.7* §2 reconciles this two-phase framing with the finer-grained burn/establishment/verification phasing used in the construction-level documents; both refer to the same process.) The commitment phase is constrained by real time and sustained device resources, making it resistant to acceleration by capital or parallelism; verification, in contrast, is constrained only by the verifier's computational resources, is local, offline-capable, and free of external coordination. This asymmetry - irreducible cost at commitment, near-zero cost at verification - is the central structural property.

The commitment phase has two components.

The first is the **collateral commitment** (the *envelope* in the canonical Lexicon §5 terminology; the artefact-meaning of *commitment* in cryptographic register): a cryptographic transcript produced through sustained verifiable work during an interval bracketed by two Bitcoin transactions. The computation consumes block hashes that are unknown until the corresponding blocks confirm, so it cannot be performed in advance and cannot be separated from its interval. Producing more commitments requires proportionally more intervals, more device-time, and more anchor transactions. Capital can buy more concurrent capacity but cannot eliminate the physical resource floor each commitment requires. The interval is defined by anchor transactions at sufficient confirmation depth; reorg handling and other operational details are specified in BACC v1.9, which also specifies the Orange Anchor instantiation.

Two mechanisms within this work are structurally load-bearing - time-binding to Bitcoin's confirmed block sequence, and memory-hard computation - and compose to produce the positive per-commitment cost floor. A third mechanism, physical-state witnessing, raises that floor further by making virtualised simulation expensive per concurrent instance. The architecture's core claim does not depend on sensor coherence being unforgeable; the detail is developed in §5.

The second is the **attribution index**: a publicly queryable record of signed observations about commitments. Events are organised into a cryptographic tree whose root is anchored to Bitcoin. Any party can publish an event; any verifier can request a proof from any operator and verify it locally against the Bitcoin-anchored root. The operator serving the data is not trusted for correctness; Bitcoin is. Interpretation is policy: a verifier may weight events by who raised them, their content, their age, or any other consideration. No authority decides what events mean. The Bitcoin-Anchored Verifiable Attribution Index (BAVAI) Reference specifies the index design.

Verification of a commitment requires only the transcript and the Bitcoin chain. Verification of an attribution event requires the event entry, its Merkle proof, and the Bitcoin-anchored root. Both are local, offline-capable, and require no trusted third party for *inclusion* (confirming a commitment or attribution event is committed against the Bitcoin chain). Full *construction verification* (confirming the bracketed work was performed faithfully) additionally requires the published Orange Anchor reference implementation, which provides the construction-verification algorithm and reference test vectors; *BAVAI Operator Specification v1.0* §7 specifies this boundary in normative form. Concretely, verification requires local transcript validation (including sensor-coherence checks where applicable), signature checks, chain-header presence proofs, and - for attribution events - Merkle inclusion proofs. It does not require re-execution of the memory-hard work performed at commitment time. Two interaction patterns make the primitive practical: in the **check-in** pattern, a holder presents a commitment as a credential to a system that recognises Orange Anchor proofs; in the **cosign** pattern, a holder links a commitment to an existing account or identity by producing a mutual signing record published to the attribution index. The *Orange Anchor Interaction Patterns v1.0* specifies both.

Three properties follow from this structure rather than from any parameter choice.

Production cost scales linearly with the number of commitments. Each commitment needs its own bracketed interval anchored to Bitcoin's block sequence. Capital buys more concurrent intervals but cannot compress any single one. Producing N commitments needs N intervals at roughly N times the per-interval cost; no mechanism in the construction reduces per-instance physical cost as N grows. Operational economies reduce the constant factor on each interval but not the per-instance floor, so total cost remains linear in N up to the operational limits of the attacker's infrastructure - capital buys more concurrent intervals, not cheaper intervals. The per-instance floor is set by physical properties of time and memory bandwidth, conditional on the memory-hard parameters being calibrated such that reduction below threshold makes commitments unverifiable; the calibration is in the SCCM and the construction paper.

Commitments are operationally non-transferable. The production transcript itself is non-transferable: it contains evidence of a specific device's hardware behaviour under sustained load during a specific interval, and no other device can reproduce that evidence without performing equivalent work in its own interval. Keys, custody, and operational control remain separately transferable, with the consequences discussed in §8; what cannot be transferred is the production context that gave the commitment its cost.

Verification needs no third party. The transcript is self-contained; the Bitcoin chain is the only external dependency and is public. Attribution events are verified the same way, against the Bitcoin-anchored root. Verification is a deterministic function of public information.

The architecture does not produce identity or enforce uniqueness. A commitment attests only that real resources were committed by a specific device during a specific interval anchored to Bitcoin. A single actor with enough resources can produce multiple commitments; the defence is the per-commitment cost incurred for each. Sybil resistance here is economic, not enumerative.

Parameter calibration, construction-specific design, proof formats, operator coordination, and protocol details are in the supporting documents and reference implementation.

5. The Collateral Commitment

A collateral commitment is produced by one device during a bracketed interval defined by two Bitcoin anchor transactions. During the interval the device performs three kinds of work: time-bound computation tied to Bitcoin's confirmed block hashes; memory-hard computation requiring sustained physical memory residency; and physical-state witnessing through the device's sensors. The cryptographic transcript of this work, plus the start and end anchors, is the commitment. Each of the three kinds of work targets a different class of adversary capability - capital-driven acceleration, parallelisation on shared hardware, and virtualised simulation respectively - and together they yield a per-commitment cost that capital, parallelism, and engineering can reduce but not eliminate.

Time-bound computation. The computation consumes block hashes from Bitcoin blocks confirmed during the interval. Because these block hashes are unknown until their respective blocks confirm, the computation cannot be performed in advance and cannot be compressed: the next block hash is not available until the next block confirms, no matter how much compute is applied. This incorporation of confirmed block hashes into the running computation is termed **reseeding** in the Lexicon (§4) and corresponds to the SCCM's *Bitcoin reseed frequency* lever (§2.2). One commitment requires one real-time interval. N commitments require N intervals, either sequentially or across N devices each running its own interval. Capital does not reduce the duration of any single interval. The cost of biasing the block-hash sequence across the calibrated interval - through withholding or grinding by a miner - is bounded below by the miner-revenue foregone in doing so, which at calibrated interval lengths exceeds any per-commitment value the bias could produce under the assumptions defined in the SCCM and BACC v1.9.

Memory-hard computation. The computation requires substantial resident memory bandwidth across the entire interval, in the tradition of memory-hard functions such as Argon2 [6]. Memory bandwidth is the primary resource limiting concurrent computation on shared hardware. Making the work memory-hard caps how many commitments a single server can produce concurrently and forces additional servers - at proportional cost - for additional concurrency. At the memory-hard parameters developed in BACC v1.9, the per-server concurrency ceiling is set by physical memory bandwidth rather than by available compute or core count, so additional concurrency scales with hardware acquisition rather than with software optimisation.

Physical-state witnessing. The device samples its own sensors throughout the interval and folds the samples into the transcript. Real devices in real environments produce sensor readings that correlate across channels as side-effects of operation: computational load shows in thermal and power sensors; motion correlates across accelerometer and gyroscope; ambient conditions correlate across light, sound, and magnetic channels. A virtualised setup must simulate every channel and keep them coherent under sustained load - a per-instance cost that grows with the number of concurrent instances.

What sensor coherence does and does not prove. Verification of physical evidence relies on statistical coherence across independent sensor channels under sustained load. A verifier accepts a transcript whose cross-channel correlations match the profile of a real device performing the prescribed work. The defence is the cost of fabricating coherent synthetic data at scale, not the unforgeability of any one signal. Sensor coherence is therefore a cost-raising mechanism, not a cryptographic proof of physical operation, and a successful coherence check should not be interpreted as one. The specific tests, channels, sampling cadence, and acceptance thresholds are specified in BACC v1.9.

The economic defence rests primarily on time-binding to Bitcoin's block sequence and memory-hard computation, which together impose a positive per-commitment cost floor that capital and parallelism cannot eliminate. Sensor coherence raises that floor further; the architecture does not depend on it being unforgeable.

Together, these three kinds of work produce a per-commitment cost that is positive, real, and resistant to amortisation. Under this construction the cost cannot be compressed below Bitcoin's block cadence, cannot be parallelised beyond the per-server memory-bandwidth ceiling, and cannot fall below the bare-metal cost of physical operation. Interval duration, memory hardness, and sensor coverage are calibration parameters developed in BACC v1.9; in particular, the architecture supports two standard burn-duration tiers - a lightweight 6-block tier and a heavyweight 144-block tier (see *Orange Anchor Lexicon v2.7*) - with the cost-positivity property holding at either tier. The structural property - positive cost, linear in volume - holds at any calibration, as established in §4.

6. The Attribution Index

Attribution is the public, permanent, verifiable record of observations made about commitment holders. A common class is the **flag** - a signed statement that a commitment has been associated with bad-faith behaviour. Flags do not seize or transfer anything; they publish a record. Because the value of a cost-backed commitment is contingent on continued legitimate recognition, observations that diminish that recognition reduce the commitment's practical value as backing for new relationships. The commitment is not confiscated; its standing is publicly altered. What any record means in any context is interpretation applied by the verifier.

Attribution events are entries in a sorted cryptographic tree, in the lineage of authenticated logs and key directories such as Certificate Transparency [7] and CONIKS [8], using sparse Merkle structures [9] for efficient inclusion and range proofs. Each entry has a key - a type indicator concatenated with the commitment's signature - and a value containing the attributing party's public key, signature over the canonical record, and a timestamp. The 32-byte tree root for each epoch is anchored to Bitcoin via a standard transaction. To check whether a commitment has attribution events, a verifier builds the expected key, queries any operator, and receives the value plus a Merkle proof. The proof is verified locally against the Bitcoin-anchored root, which the verifier reads directly from the chain. An operator cannot forge a proof against a root they did not produce. The reference index construction provides authenticated range-query semantics sufficient to detect omission of entries within queried key ranges; details are in the BAVAI Reference. Different operators may serve the same data with different availability and pricing; verification produces the same result regardless of source.

The index is not limited to flags. Any signed observation that can be expressed as a key/value pair - endorsements, scores, credential references, cosigning records, attestations of any kind - can be recorded in the same structure under a different type indicator. The verification model is identical: anyone may publish, anyone may query, and interpretation is per-verifier policy. The BAVAI Reference specifies key encoding, tree structure, and operator coordination.

Several properties follow from the design rather than from any policy choice. Operators may withhold data; verifiers needing high availability can query multiple operators or maintain their own copies. Because anyone may publish, spam, collusive flagging, and adversarial publication are possible; filtering is applied by verifier policy, not at the index layer. Events published in the clear, including cosigning records, create publicly observable linkages between commitments and external identifiers; participants who require unlinkability must avoid such publication or use constructions outside this architecture. The cadence at which roots are anchored to Bitcoin sets a practical upper bound on attribution freshness. A natural verifier policy weights attribution events by the production cost of the publishing commitment, subjecting adversarial flagging to the same cost-linearity that defends commitments themselves.

7. Verification and Interaction

Verifying a commitment requires two things: the transcript the holder provides and access to the Bitcoin chain to confirm that the start and end anchors sit at the claimed block positions. Both are public. No issuer is contacted; no platform is queried; no authority participates in *inclusion verification* against the Bitcoin chain. *Construction verification* - verifying that the bracketed work was performed faithfully - is performed using the published Orange Anchor reference implementation against the verification package; the reference implementation is the open-source artefact that supplies the construction-verification algorithm and reference test vectors. Verification is deterministic from public information and can be performed offline once chain headers and the reference implementation are available locally. Attribution-event verification follows the same pattern, against the Bitcoin-anchored root.

Existing verification typically requires real-time communication with a trusted authority - an identity provider, credential issuer, or session manager - at verification time. Commitment verification removes that requirement, so verification of cost-backed claims is structurally independent of any third party's availability. This independence is scoped to the operations the architecture covers; revocation, freshness, authorisation, and other functions of existing identity infrastructure are orthogonal and not addressed here.

In the **check-in** pattern, a holder presents a commitment as a credential. The system issues a challenge - typically a random value combined with the system's identifier and a timestamp - and the holder signs it with the key associated with their commitment. The system verifies the signature against the commitment's public key and verifies the commitment itself against Bitcoin. No identity provider participates. The system applies its own policy to the commitment, any prior history it holds about the holder, and the public attribution record.

In the **cosign** pattern, a holder links a commitment to an existing account, credential, or identity. The holder signs a canonical linking message twice - once with the commitment key, once with the key of the external identity - and publishes the pair as an attribution event. Any verifier later querying either side can find the linkage and verify it. This is how commitments compose with existing identity systems - decentralised identifiers, verifiable credentials, social-network keys (e.g. Nostr npubs), and other frameworks whose keys can sign canonical messages. The commitment imports cost-backing into those systems; those systems import their recognition into the commitment.

8. Economic Architecture

For a single-device producer, the cost of a commitment is the cost of running a background process on hardware they already own, during otherwise-idle time. The dominant measurable cost is the Bitcoin anchor fees for the start and end transactions, and these can be amortised by batching operators that aggregate many commitments into fewer Bitcoin transactions. Device time and energy are marginal relative to existing use for the honest reference class defined in the SCCM. The cost is paid once and the same commitment backs interactions across any number of systems over its useful life.

For a concurrent producer producing many commitments at scale, the cost structure resists amortisation at every layer. The time-bound computation cannot be compressed - each commitment needs its own real-time interval. The memory-hard computation cannot be parallelised beyond the per-server memory-bandwidth ceiling. Physical-state witnessing cannot be shared across instances - each must generate its own sensor profile. Per-commitment physical cost stays positive and scales linearly with volume (as established in §4). Bulk procurement, energy contracts, and fee batching reduce per-interval cost but do not eliminate the per-instance physical floor.

Bitcoin footprint and batching model. Each commitment requires two anchor points, which may be realised through batched start and end transactions. At typical fee levels this dominates user-facing cost; batching operators aggregate many commitments into single Bitcoin transactions, amortising fees while introducing a semi-trusted intermediary layer. Under typical batch sizes the marginal Bitcoin footprint per commitment is small relative to a standard transaction, comparable to other established uses of the commitment surface for embedding application data. The model assumes operators are incentivised by fees and reputation; users may choose multiple operators to mitigate censorship or delay. Full trust assumptions, failure modes, and fee economics are detailed in BACC v1.9.

At this positive per-commitment cost floor, the economic model of synthetic participation at scale changes structurally. Today's synthetic-account economy depends on near-zero marginal cost per object; a positive floor removes that condition. The full cost analysis, including the most favourable assumptions for sophisticated concurrent producers and specific operational adversary setups, is in the Strategic Cost Calibration Model and the companion Adversarial Scenario Analysis. The SCCM defines an adversary with substantial capital, virtualisation tooling, and operational sophistication, and specifies the methodology by which the per-commitment cost floor against that adversary is composed from a finite set of structural levers. The dominant user-facing cost component under typical operation is amortised Bitcoin anchor fees under batching.[1] The architectural claim in this paper is the structural one: per-commitment cost is positive and scales linearly with volume under the stated threat model. Magnitude is calibration; positive linear scaling is structural.

The attribution index reinforces the economic structure. A commitment used in bad faith may receive attribution events that any verifier can discover. The cost was sunk at production; if attribution accumulates, the commitment's value as backing for new relationships diminishes without anyone recovering the original cost. The producer's calculation is the production cost weighed against the value they can extract before attribution accumulates. Earlier attribution worsens the calculation for bad-faith

producers; honest use carries the same cost but accumulates positive standing. This reinforcement assumes verifiers query the index - where index availability is restricted, the marginal deterrent is reduced. The production-cost defence does not depend on attribution; the reinforcement does.

The four elements - production cost, attribution record, recognition value, and verification independence - are mutually reinforcing. Cost makes recognition real. Recognition makes verification useful. Verification independence lets recognition compose across systems. Attribution keeps the cost binding under bad-faith use. None is meaningful in isolation; together they are self-reinforcing under honest use and self-limiting under bad-faith use.

The economic argument rests on a stated threat model: an adversary with substantial capital, computation, virtualisation tooling, and operational sophistication, attempting to produce commitments at scale. Against such an adversary the cost-linearity property holds - capital buys more concurrent intervals but does not reduce per-interval cost.

Several adversary capabilities lie outside this defence and are acknowledged explicitly. A physically compromised device can produce a commitment under an attacker's control: the work is real, but the producer of record may not be the device's nominal owner. Sensor evidence can in principle be spoofed by an attacker with sufficient physical control; coherence raises the cost but does not eliminate it. Custody of the private key can be transferred; the commitment remains bound to its production context, but the party operating the key may change. These vectors affect the binding between one commitment and one party, not the linearity of cost across many commitments.

Secondary markets for aged or transferred commitments are possible. A buyer may acquire an existing commitment rather than produce one. Such transfers do not reduce the original production cost, and, crucially, the supply of commitments on any secondary market is bounded by the volume actually produced historically - a buyer cannot conjure one, only acquire one that already exists. The economic defence therefore remains linear in the number of commitments in existence, regardless of how custody is subsequently distributed. Attribution history travels with the commitment, so a buyer inherits whatever standing it carries; aged commitments with clean history command a premium, and that premium is the production cost expressed through the market.

8.1 Threat Model Summary

The table consolidates the threat model that the preceding sections develop narratively. It is the basis on which the cost-linearity claim is asserted, and the basis on which the limitations in §10 are scoped.

Adversary capability	Scope	Defence or acknowledged limit	Reference
Pre-computing commitments before their interval	In-scope	Block hashes consumed during the interval are unknown until the interval's blocks confirm	§5 (time-bound)
Parallelising production on shared hardware	In-scope	Memory-hard computation caps per-server concurrency; additional concurrency requires additional servers	§5 (memory-hard), [6]
Simulating physical operation in virtualisation	In-scope	Cross-channel sensor coherence under sustained load raises the per-instance simulation cost	§5 (physical-state)
Capital scaling to produce many commitments	In-scope	Per-commitment cost is positive and scales linearly with volume up to operational constants	§4, §8
Operator withholding or forging attribution proofs	In-scope	Roots are Bitcoin-anchored; authenticated range-query semantics detect omissions; verifiers may query multiple operators	§6, [7], [8], [9]
Spam, collusive flagging, adversarial publication to the index	In-scope (at policy layer)	Anyone may publish; filtering is verifier policy, not an index-layer guarantee	§6
Physical compromise of a device producing a commitment	Acknowledged limit	The work is genuine, but the producer of record may not be the device's nominal owner	§8, §10
Sensor spoofing under sufficient physical control	Acknowledged limit	Coherence raises but does not eliminate spoofing cost; not cryptographic proof of physical operation	§5, §10
Key custody transfer or sale of an existing commitment	Acknowledged limit	Production transcript is non-transferable; keys and custody are separately transferable; secondary-market supply is bounded by historical production	§4, §8, §10

The in-scope adversary is the one against which cost-linearity is claimed. The acknowledged limits affect the binding between one commitment and one party; they do not affect the linearity of cost across many commitments. The out-of-scope items are not failures of the architecture but boundaries of its claims.

8.2 Structural Consequences

Six consequences follow from the structural properties established above. They are not parallel; each amplifies the next. None is a calibration parameter.

Cost-curve shape. This changes the cost structure of synthetic participation from nearly flat to linearly proportional to N. Any production model whose viability depends on near-zero marginal cost per object is structurally incompatible with the construction.

Capital neutralisation. Capital expenditure by a concurrent producer scales the number of commitments produced, not the per-commitment cost. The construction provides no mechanism by which additional capital compresses per-instance resource expenditure.

Compounding standing under honest use. A commitment used honestly across multiple systems accumulates attribution standing over time. The production cost is paid once; the standing it enables accumulates with continued honest use.

Verification permanence. Verification is a deterministic function of public information. A commitment verified today against the Bitcoin chain remains verifiable in the same way in ten years, without any ongoing coordination with an issuer or operator.

Cross-system amortisation. A single commitment, once produced, is usable as a cost-backed claim across every verifying system that accepts the Orange Anchor proof format. The production cost is amortised across the full set of systems a holder interacts with over the commitment's useful life.

Asymmetry durable across hardware generations. The cost asymmetry depends on relationships between resources - time-binding versus capital, memory bandwidth versus concurrency, per-instance witnessing versus virtualisation - not on absolute magnitudes of any one resource. Hardware improvement scales honest and adversary capability symmetrically; the structural asymmetry is preserved.

The compound effect: honest participants pay production cost once and accumulate standing across systems over time; concurrent adversaries pay per-commitment cost for each fake commitment and gain nothing recoverable. The asymmetry is not just instantaneous - it widens with adoption and time.

9. Implications

When verification needs no real-time coordination, a class of costs that exists only because of coordination becomes structurally unnecessary. Authentication infrastructure, identity providers, certificate authorities, and trust intermediaries exist in part because trust is currently coordinated at verification time. For uses where commitment verification replaces coordination-based verification, that requirement is gone. Systems built on this foundation are structurally independent of third-party coordination for the verification operations the architecture covers.

Consider a forum that wants to limit synthetic accounts without imposing identity verification. Today it must accept whatever volume automation produces, impose CAPTCHA-style friction that costs legitimate users without meaningfully costing automation at scale, or require third-party identity verification that introduces dependency and exclusion. Under this architecture, the forum accepts Orange Anchor commitments as a sufficient credential. A legitimate user produces one commitment on their own device and uses it across any forum that recognises the same primitive (where cross-context linkability is acceptable). A producer running many synthetic accounts pays per-account production cost that scales linearly with the number of accounts. The forum obtains Sybil resistance without depending on a third-party identity provider or trusted verification service at verification time; the commitment can additionally be cosigned to the forum's internal account identifier to produce a verifiable linkage. The structural cost of the defence is borne by the synthetic-account producer, not by the platform or the legitimate user.

One possible application is commitment-weighted voting in governance systems. Weighted voting rights can be tied to cost-backed commitments rather than to token holdings or to identity proofs. Token-weighted voting is subject to capital accumulation; identity-weighted voting is subject to issuer gatekeeping. Commitment-weighted voting substitutes per-credential production cost for both (without claiming one-person-one-vote semantics).

The architecture composes with existing identity systems rather than replacing them. A commitment can be cosigned to a decentralised identifier (DID), a verifiable credential, a social-network key, or any other system whose keys can sign canonical messages. The commitment imports verifiable cost; those systems retain their existing semantics and gain a new verifiable attribute. Adoption requires no internal change to those systems.

The same pattern - off-chain construction of complex structures, on-chain commitment via Bitcoin, local verification against the commitment surface - applies beyond identity. Supply-chain provenance, AI content attribution, governance, reputation, and any domain that needs verifiable claims grounded in cost can adopt the same pattern with different specific constructions. Orange Anchor is one instance of a more general architectural class. The general class is the durable contribution; specific constructions will evolve.

The calibration target for any real deployment is a window: high enough to create meaningful economic friction against synthetic participation, yet low enough that the cost remains practical for legitimate users. The specific bounds of that window are developed in the Strategic Cost Calibration Model.

10. Limitations and Non-Goals

The architecture is deliberately scoped as an economic primitive for cost-backed claims. The following limitations are stated explicitly because the architectural argument is strongest when its scope is honest - and because each limitation is a boundary of the claim, not a failure of the architecture.

- It does not establish or verify legal identity, biometric identity, personhood, or uniqueness - such proofs require composition with other attestations.
- It does not prevent all Sybil attacks; it changes their economics by imposing a positive per-commitment cost floor. Sufficiently capitalised attacks remain possible but become economically irrational at meaningful scale.
- It does not guarantee that every commitment is produced by its nominal owner (device compromise, key custody transfer, and operation on someone else's behalf are possible).
- It does not provide cryptographic proof of physical operation - sensor coherence raises fabrication cost but does not constitute unforgeability (see §5).
- It does not prevent secondary markets in aged or transferred commitments; supply on such markets is bounded by historical production volume.
- It does not provide unlinkability across contexts: a commitment used across multiple systems becomes a stable identifier across them. Participants requiring unlinkability must produce separate commitments per context, at separate production cost.
- It does not replace existing identity, credential, authorisation, or trust infrastructure; it is a composable primitive.

These limitations are stated explicitly because the goal is a defensible economic primitive with bounded, verifiable properties - not a general-purpose identity or trust solution. The architecture is strongest when its scope is honest.

10a. Scope and Forthcoming Work

This paper, the *Bitcoin-Anchored Collateral Commitments* construction paper, the *Orange Anchor Lexicon v2.7*, the *BAVAI Reference*, the *BAVAI Operator Specification*, the *Orange Anchor Interaction Patterns*, the *Strategic Cost Calibration Model*, and the *Orange Anchor Integration Brief* constitute the published Orange Anchor document suite at this release. Architecture and protocol surfaces in the released suite are stable for review and integrator implementation. The following supplementary documents are scheduled for subsequent release; calibration is treated as continuing discipline rather than a one-time setup:

- **Calibration Annex** - will publish preliminary numerical κ ranges for the Strategic Cost Calibration Model under explicit threat-model assumptions; the SCCM §5 weights remain calibration-pending until this document is published.
- **Adversarial Scenario Analysis** - will develop concrete attack scenarios (Sybil farms at varying capital levels, virtualisation attacks at varying sophistication, anchor-front-running, attribution-spam) against the calibrated parameters from the Annex.
- **Integrator Quick Start** - will provide an end-to-end walk-through of integrating Orange Anchor verification into a verifier application, complementing the *Integration Brief*'s architectural framing with concrete code paths against the reference implementation.
- **Orange Anchor Technical Paper** - will provide formal cryptographic analysis of the construction's security properties and parameter-calibration justification, complementing the architectural argument made here.
- **Reference Implementation** - the open-source codebase that supplies the construction verification algorithm, reference test vectors, and `proof_internals` slot layout deferred from *BAVAI Operator Specification v1.0* §7.4.

Readers evaluating Orange Anchor for production deployment should treat construction verification as available only once the reference implementation is published; the protocol surfaces specified in the released suite remain stable across that transition.

11. Conclusion

This paper has proposed an architecture in which real resource expenditure, anchored to Bitcoin's public block sequence, backs a self-issued digital claim. The collateral commitment is produced through sustained verifiable work on a specific device during a defined interval, anchored to Bitcoin and verifiable by anyone with the chain. The attribution index records public observations as peer-to-peer entries that any verifier can query and interpret by their own policy. The two interaction patterns - check-in and cosign - make the architecture usable both for systems recognising commitments and for holders linking commitments to existing identities. The economic structure produces positive cost on production that scales linearly with volume up to operational constants, preserves that cost as attributable scarcity across every system that recognises the commitment, and keeps verification free and offline-capable. The asymmetry undercuts the economic basis for synthetic participation at scale under the stated threat model.

Supporting documents accompany this paper: the *Bitcoin-Anchored Collateral Commitments* construction paper (commitment pattern, calibration, Orange Anchor instantiation); the *Orange Anchor Lexicon v2.7* (canonical vocabulary, envelope/burn/burn-tier definitions, phase-framing reconciliation); the BAVAI Reference (attribution index, key structure, operator coordination); *Orange Anchor Interaction Patterns v1.0* (check-in and cosign in integration detail); the Strategic Cost Calibration Model and Adversarial Scenario Analysis (economic calibration under aggressive attacker assumptions); and the reference implementation. The architecture is proposed; the construction is one instance among possible instances; the pattern is what should survive scrutiny. The invitation is to examine, critique, build, and extend.

References

[1] J. R. Douceur. *The Sybil Attack*. In Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS), 2002.

[2] A. Back. *Hashcash - A Denial of Service Counter-Measure*. 2002.

[3] M. Borge, E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, and B. Ford. *Proof-of-Personhood: Redemocratizing Permissionless Cryptocurrencies*. In IEEE Security and Privacy Workshops, 2017.

[4] Worldcoin Foundation. *Worldcoin: A New Identity and Financial Network*. 2023.

[5] S. Nakamoto. *Bitcoin: A Peer-to-Peer Electronic Cash System*. 2008.

[6] A. Biryukov, D. Dinu, and D. Khovratovich. *Argon2: New Generation of Memory-Hard Functions for Password Hashing and Other Applications*. IEEE European Symposium on Security and Privacy, 2016.

[7] B. Laurie, A. Langley, and E. Kasper. *Certificate Transparency*. RFC 6962, IETF, 2013.

[8] M. S. Melara, A. Blankstein, J. Bonneau, E. W. Felten, and M. J. Freedman. *CONIKS: Bringing Key Transparency to End Users*. In USENIX Security Symposium, 2015.

[9] R. Dahlberg, T. Pulls, and R. Peeters. *Efficient Sparse Merkle Trees: Caching Strategies and Secure (Non-)Membership Proofs*. In Nordic Conference on Secure IT Systems (NordSec), 2016.

---

contact@bacc.at

DRAFT v2.4 - 15-05-2026

---

1. The calibrated magnitude of the floor is developed in the Strategic Cost Calibration Model. Because the dominant user-facing component is amortised Bitcoin anchor fees, the magnitude is sensitive to the prevailing fee market. Specific figures, sensitivity analysis, and the full adversary specification belong in that document rather than in the architectural claim made here. ↩