Cascade: Citation-Driven Value Propagation for Creative Works

Adam Carlson

June 2026

1 Abstract

Cascade is a citation-driven value-propagation architecture for creative works. Where existing intellectual-property regimes treat influence as a liability to be policed, and where existing royalty mechanisms depend on adversarial enforcement that has proven structurally fragile, Cascade makes citation itself the mechanism by which economic value flows from new works upstream to all the works that influenced them. A single act of patronage at a downstream work cascades through generations of citation lineage in one transaction, compensating ancestral creators across the network without requiring enforcement, agreement, or any external resolution event. The architecture is self-contained — value is determined by network position rather than by external markets — and is designed to augment existing intellectual-property regimes rather than replace them. This whitepaper describes the architecture, its mechanics, and its structural properties, and explains why citation-as-economic-mechanism produces a creative economy with incentive alignments unlike any existing model.

2 The Problem: Influence as Liability

Intellectual property law operates on an adversarial premise. When a creative work resembles or borrows from an earlier one, the legal system treats the resemblance as a liability for the new creator and attempts to compensate the earlier creator through enforcement actions, licensing requirements, or damages. Influence — the central mechanism by which culture and knowledge advance — is structurally treated as a problem to be policed rather than a value to be channeled.

Attempts to address this asymmetry have produced structurally fragile alternatives.

NFT royalties promised ongoing compensation to creators on resale, but the regime collapsed in practice because enforcement required cooperation from secondary marketplaces and settlement systems that ultimately treated royalty enforcement as optional. The mechanism worked only when downstream parties chose to honor it; the moment marketplaces decided not to, the entire compensation flow stopped.

Statutory royalty schemes depend on creators having the resources to monitor downstream use and pursue enforcement. Most creators don't. The mechanism exists on paper and rarely fires in practice.

Academic citation systems track intellectual lineage faithfully — a foundational paper that becomes the basis for an entire field is widely cited, the citations are visible, the influence is undisputed — but no economic component is attached. The original author receives reputational benefit through external institutions and nothing else.

Patent licensing requires adversarial negotiation, costly enforcement, and active monitoring of infringement. The mechanism rewards the lawyers more reliably than the inventors.

Each of these approaches makes the same structural assumption: that intellectual lineage and economic value must be tracked, enforced, or settled through separate mechanisms. Provenance is one system. Compensation is another. The mechanism that records *who built on whom* is divorced from the mechanism that flows *value among them.* The result is that the act of acknowledging influence — the act of citing — generates no value flow on its own. It generates value only when downstream parties cooperate to enforce a separate compensation mechanism, and they routinely don't.

What's missing is an architecture in which acknowledgment of influence *is itself* the act that routes value. Where provenance and compensation are unified rather than separated. Where citation isn't just a tracking mark — it's the wire.

3 The Insight: Citation as the Wire

What if the act of citing a prior work — the act of acknowledging influence — were itself the act that routes economic value? Not as a parallel mechanism that requires separate enforcement, not as a tracking system that has to wait for downstream parties to honor it, but as the *primary* mechanism by which value flows through the network?

This is the architectural insight at the heart of Cascade. Citation becomes load-bearing. When a new work registers and identifies the prior works that influenced it, that identification is not merely metadata. It is the wire through which subsequent patronage of the new work will flow back to the works that made it possible.

Three things follow from this inversion.

First, provenance and compensation are unified rather than separated. The mechanism that records who built on whom is the same mechanism that flows value among them. A single act — citation — does both.

Second, no enforcement is required at the boundary. The downstream parties do not have to cooperate with a separate royalty system, marketplace, or settlement layer. The flow of value is intrinsic to the citation graph itself, computed by the network's own rules and routed by the network's own mechanisms. There is nothing external to enforce, because the wire is internal.

Third, citation honesty becomes economically incentivized. Under-citation is no longer a cost-free omission with reputational consequences that may or may not materialize. Citation directly allocates economic stake in the new work to the prior work's stakeholders, which means the new creator's incentive to cite honestly aligns with the prior creator's interest in being cited. The two halves of an honest acknowledgment now point in the same direction.

The remainder of this whitepaper describes the architecture that makes this work — how the wire is constructed, how value enters it, how it propagates upstream through generations of influence, and what structural properties emerge from a network organized around influence-as-economic-flow.

4 The Architecture

The architecture has six interacting components: a citation graph, origin shares, patron shares, the cascading propagation mechanism, a treasury, and a ghost-node mechanism for off-chain prior works. Citation curation — the determination of what counts as influence — operates as a layer above the architecture and supports multiple governance modes. This section describes each component in the order they operate.

4.1 The Citation Graph

The citation graph is a directed acyclic graph maintained on a distributed ledger. Nodes represent registered works. Edges represent citations from a citing work to a cited prior work. Each node carries an origin-share register tracking who holds stake in the work and a patron-share register tracking outstanding shares purchased by community participants. The ledger additionally records share holdings, transfers, and proceeds distributions.

The architecture is not tied to any particular ledger technology. Implementations may use public blockchains supporting programmable smart contracts, permissioned ledgers operated by federations of platforms, or hybrid configurations where the graph is maintained off-chain with cryptographic commitments stored on-chain.

4.2 Origin Shares

Origin shares represent stake in a registered work. When a new work is registered with a set of citations, the system allocates the work's origin-share pool among three categories: a *creator allocation* reserved for the creator of the new work; a *citation allocation* distributed across the cited prior works; and a *treasury allocation* directed to a community-controlled treasury account (described in Section 4.5).

The split is governed by a configurable allocation function. The function may be a fixed split or may respond to characteristics of the work — the number of cited prior works, the work's classification, the state of the treasury, citation weights, or other factors. Allocation parameters are typically set by governance and modifiable through a governance procedure.

Origin shares allocated to cited prior works are credited to the citation graph rather than transferred directly. The cited work's existing origin-share holders become indirect holders of the new work's origin shares through the cascading propagation mechanism. This indirection is what makes the cascade possible.

4.3 Patron Shares and Patronage Purchases

Origin shares represent who is entitled to value from a work. Patron shares are the mechanism by which value enters the system in the first place. A new work makes patron shares available for community purchase. The supply, pricing curve, and availability schedule are configurable per work — set at registration, determined by a bonding curve, governed by an automated market maker, or any combination.

When a community participant purchases patron shares, the patronage module records the purchase, transfers the patron shares to the buyer, and routes the proceeds of the purchase to the propagation mechanism described next. Patron shares confer entitlement to upstream revenue distributions but do not confer ownership of the underlying work or its rights under existing intellectual-property regimes. They are transferable on secondary markets, but the cascade propagation flow is triggered only by primary issuance — buying a patron share is what enters value into the network; reselling that patron share later does not.

4.4 The Cascade — Upstream Value Propagation

This is the architecture's defining mechanism. When a patron buys shares of a work, proceeds do not flow only to the work's creator. They cascade through the entire upstream citation lineage of the patronized work, compensating ancestral creators across multiple generations in a single transaction.

The mechanism is recursive. Proceeds entering a work are distributed across that work's origin shareholders proportional to their holdings. Because a portion of those origin shares is held by upstream cited works, the portion allocated to each cited work is then further distributed — through *that* work's origin shareholders, including the works *it* cited — and so on, walking the citation graph backward through every chain of influence.

A worked example. Consider works W1, W2, W3, and W4, where W4 cites W3, W3 cites W2, and W2 cites W1. A patron purchases shares of W4. The proceeds split: a portion to W4's creator (the creator allocation); a portion routed to W3 (the citation allocation); a portion to the treasury. The portion routed to W3 then splits: a portion to W3's creator, a portion routed to W2. The portion routed to W2 splits: a portion to W2's creator, a portion to W1. The portion at W1 goes to W1's creator, since W1 has no further citations. One patron transaction at W4 compensates four creators across three generations of citation lineage — without requiring any of those creators to take action, without enforcement, and without dependence on any external resolution event.

The cascade is configurable. A *depth parameter* limits how many upstream hops a single purchase reaches, with deeper propagation rewarding distant ancestry and shallower propagation concentrating value in proximal influences. A *per-hop decay parameter* reduces the proportion of value reaching each successive upstream hop, providing a tunable mechanism for shaping how value spreads across the depth of the ancestral tree. These parameters can be set globally, per-work, or per-category by governance.

In other words, citation is the wire through which patronage of a downstream work compensates upstream creators, and the citation graph is the wiring diagram. The cascade walks the wiring automatically. No one has to ask. No one has to enforce. No one has to know. When a patron buys shares of a 2026 work that cites a 2020 work that cited a 2015 work, the 2015 creator earns from the 2026 patronage as a structural consequence of the citation relationships, not as the result of any agreement between them.

4.5 The Treasury

The architecture supports a community-controlled treasury that participates as a structural party in every registration. A configurable percentage of every work's origin-share pool — the *treasury allocation* — is credited to a treasury account at registration time, in the same manner that allocations to the creator and to cited prior works are computed. The treasury's holdings consist of origin shares of registered works, and the downstream entitlements those holdings accumulate as the cascade routes value upstream from later patronage.

The treasury is a productive economic agent of the community that operates it, not merely an accumulator. Treasury disbursements are authorized through whichever governance procedure the implementation uses — decentralized autonomous organization voting, multi-signature trustee approval, foundation governance, or algorithmic disbursement under predefined rules — and may fund grants to creators in designated categories, subsidies for ghost-node conversion costs, operating reserves for infrastructure, bounties for security audits and AI component evaluation, matching contributions to creator-led patronage campaigns, or other activities authorized by governance.

In some implementations, the treasury is configured to participate as a patron in the patronage mechanism itself — purchasing patron shares of works the community has identified as deserving of support, becoming an upstream entitlement holder in those works, and producing a reflexive mechanism by which treasury value flows back into the network rather than only out of it.

In other words: a fraction of every new work flows to a shared pool, the community decides how the pool is used, the pool grows or contracts with the network's overall activity, and the treasury may itself act as a patron — becoming an upstream creator-of-record in works it supports. The treasury is structurally a configurable economic instrument whose specific applications are determined by the community that operates it.

4.6 Ghost Nodes — Off-Chain Prior Works

The citation graph would misrepresent influence relationships if it terminated at the boundary of on-chain works. Most creative lineage predates any system that might host it, and truncating the graph at the on-chain boundary would create economic pressure to under-cite older works on the basis that they cannot claim their shares. The architecture resolves this through *ghost nodes*: placeholder nodes in the citation graph that represent prior works existing outside the on-chain ledger.

When a creator cites an off-chain prior work, the registration module creates a ghost node as a placeholder. The ghost node participates structurally in the citation graph in the same manner as a regular node — it appears as an upstream citation source, it receives an allocation of origin shares of the citing work, and it accumulates downstream entitlements from subsequent patron purchases that cascade through the citing work. The difference is that the ghost node's accumulated entitlements are held in an escrow account rather than distributed to a creator-of-record.

The system provides a conversion mechanism. An off-chain creator may convert their ghost node into a regular node by submitting a transaction that establishes their on-chain identity, demonstrates their claim to the off-chain work, **and identifies the citations of the newly-claimed work — the prior works that influenced it, which may themselves include further off-chain works represented as new ghost nodes.** Upon conversion, **the accumulated escrow is treated as if it were patron-purchase proceeds entering at the newly-claimed node: a portion flows to the converting creator, the citation portion cascades upstream through the newly-declared citations and pays down through whatever lineage the off-chain work itself rests on, and a portion goes to the treasury.** The converted node thereafter participates in normal upstream propagation.

This produces a structurally recursive property. Each conversion event has the potential to extend the off-chain lineage further back in time, and each newly-discovered ancestral ghost node accumulates escrow of its own that motivates yet another conversion. The graph back-fills as off-chain creators surface to claim their place — each claim potentially summoning more.

The release mechanism makes participation self-incentivizing. An off-chain creator becomes economically motivated to register not because of any abstract appeal but because real accumulated earnings — generated by downstream citations and propagated patron purchases — are awaiting claim. The system periodically publishes per-ghost-node escrow balances to encourage discovery and conversion. Ghost nodes are themselves citable by later works, producing multi-level escrow accumulation across off-chain lineage. Safeguards against abuse include public review periods for ghost-node citations, community governance procedures for contested attributions, and conversion fees that recoup maintenance costs from accumulated escrow at the time of claim.

In other words: the citation graph models the actual creative-influence network, including works that predate the system or live outside it. On-chain participation is the mechanism for claiming one's place in the network, not the prerequisite for being represented in it — and each act of claiming may reach further back, summoning more of the lineage into representation with it.

4.7 Citation Curation

Sections 4.1 through 4.6 describe what happens once citations are determined. The behavior of the entire architecture depends critically on the input that drives the cascade: *what counts as a citation*. The architecture supports three citation-curation modes, configurable at the work level, the ledger level, or by community governance.

Creator-direct submission. The creator names the citations themselves at registration time. Authority sits with the creator. This mode is suitable when the creator has comprehensive knowledge of the influences and can be relied upon to disclose them honestly. Cascade's economic alignment helps here — under-citing prior works directly reduces the citation allocation away from creators with legitimate claims to influence stake, which means the upstream creators have a real economic interest in being cited, which produces incentive pressure toward honest disclosure even without enforcement.

Community proposal and voting. Community participants propose citations linking a new work to specific prior works, and each proposed citation is subjected to a voting procedure. Citations that meet a configurable approval threshold are incorporated; those that fail are excluded. Voting weight can be uniform per participant, proportional to share holdings, or computed by a governance-defined function combining multiple factors. This mode places authority over citation determination with the collective rather than the individual creator and is well-suited to contexts where the network has a strong interest in independently determining influence.

AI proposal with lazy consensus. An artificial-intelligence component analyzes the new work and proposes a set of citations. The proposed citations are published in a public-review queue for a configurable review period; absent objections exceeding a configurable threshold during the review period, the proposed citations are automatically incorporated. Inaction during the review period is treated as approval. This mode produces throughput that per-citation voting cannot achieve while preserving the community's ability to override any specific proposal.

The AI-proposal mechanism operates under *two-level governance* that structurally separates the rules from the cases. At the *policy level*, community governance votes on the rules under which the AI operates — what algorithms identify candidate citations, what thresholds qualify a candidate, what categories the AI is or isn't authorized to propose for, how the AI is retrained and audited. At the *case level*, the AI executes within the policy-level rules without per-citation voting; the lazy-consensus review period provides override for individual cases. Democratic authority sits where deliberation is feasible (the rules); automated execution operates where deliberation does not scale (the cases).

The three modes can operate in parallel for the same work, with citations produced through any mechanism incorporated into the citation set subject to the validation procedures of that mechanism. In other words: *who decides what counts as influence is a design parameter rather than a fixed assumption.* The system supports creator judgment, collective judgment, and AI judgment under community-set rules — three distributions of authority over the citation graph, each appropriate to different contexts, and each configurable per work.

5 Cascade as a Knowledge Graph

The citation network described in Section 4 is more than a routing structure for value. It is a *typed knowledge graph* in the formal sense — a substrate in which nodes carry meaning, edges encode specific intellectual relationships, and the structure of the graph itself becomes legible to systems that reason over it. Recognizing this opens the architecture to a broader class of operations than value propagation alone, and connects it structurally to a separate body of work on Reflexive Knowledge Graphs as a paradigm for systems that build and reason over their own substrate.

Two structural properties make this framing natural.

5.1 The Citation Network as a Typed Graph

Cascade's citation graph has nodes (works) and directed edges (citations) — but the edges are not undifferentiated pointers. Each citation carries a type and may carry weight and rationale. A citation from a new work to a prior work asserts a specific intellectual relationship: *this work draws on*, *this work builds upon*, *this work extends*, *this work responds to.* The type schema for citation edges is configurable per domain — what counts as a meaningful relationship between a research paper and its predecessors differs from what counts between a software library and its dependencies — but in every implementation, the typing is what makes the graph *legible* rather than merely *navigable.* Algorithms that reason about the structure of the network can read typed edge sequences as argument chains, identify clusters of mutually-relating works, and detect topological features that imply meaning beyond what any individual node states. The cascade itself is the simplest operation that exploits this structure; many more are available.

5.2 Each Work as a Sub-Graph

Beyond the network-level graph of works-and-citations, each individual work can itself be modeled as a typed knowledge graph of its internal structure. A song registered in Cascade can be associated with a sub-graph of its components — verses, hooks, samples, instrumental parts — connected by typed relationships describing how those components contribute to the work. A research paper can be associated with a sub-graph of its claims, the hypotheses they rest on, the evidence they cite, and the syntheses they produce. A software library can be associated with a sub-graph of its modules, interfaces, dependencies, and tested behaviors. The type schema for inner-work graphs is domain-specific, and different communities operating different ledgers may adopt different schemas appropriate to the kinds of works they host.

This inner-graph structure does not change how the cascade routes value at the outer-graph level — the cascade still propagates through citation edges between works. What inner-graph structure enables is *finer-grained citation*: a new work can cite specific components of a prior work rather than the work as a whole, and influence relationships can be expressed at the level at which they actually operate. A software library can cite the specific data structure from another library that it adopted; a paper can cite the specific argument from another paper that it extends; a song can cite the specific sample from another song that it builds on. Inner-graph citation makes the network's representation of influence more precise.

5.3 The Reflexive Property

Cascade's architecture is reflexive in a sense familiar from the broader Reflexive Knowledge Graph paradigm: the same system that constructs the substrate also operates on what it has constructed. New registrations build the graph. The cascade reads the graph it has built to route value. Citation-curation modes — particularly AI proposal — read the graph and propose additions to it. Topology-aware operations described in Section 7 read the graph and surface structural features for human evaluation. Each operation acts on the structure produced by every prior operation, and produces structure that subsequent operations will act on.

This reflexivity is what makes Cascade more than a payment routing system. The graph is not a side effect of the value propagation; the graph is the substrate. Value propagation is one of several operations the substrate supports. Other operations — discovery of latent intersections, identification of structural gaps, surfacing of works whose network position implies disproportionate influence — operate on the same substrate and produce different kinds of output. Section 7 develops one such operation in detail.

6 Properties

The architecture produces a set of structural properties worth naming. Each is a structural consequence of how the architecture is built rather than an external constraint layered on after the fact. They emerge from the same mechanics — and naming them clarifies what makes Cascade structurally different from prior systems for compensating creative work.

6.1 Self-Contained Economy

Cascade is a self-contained economic system. Every input to it — work registration, citation, patron purchase, treasury disbursement, ghost-node conversion — originates within the system itself. Every value flow is computed and routed by the system's own rules. No external valuation signal, marketplace, or commercial performance metric participates in determining what anything in the system is worth. The economic state at any moment is fully described by the citation graph, the share registers, the treasury holdings, and the ghost-node escrow balances. A patron purchasing patron shares is investing in the work's structural position within the network, not in any externally-resolved performance event, and the patron's return depends entirely on subsequent activity within the system.

This distinguishes Cascade from prior valuation systems whose routing functions depend on signals outside themselves. In equity markets, share price reflects external business performance. In secondary-marketplace royalty systems, value flows depend on resales on external platforms. In subscription systems, value enters through external payment channels and the system acts only as routing. In citation-only academic systems, intellectual lineage is tracked but no economic mechanism is associated with the tracking. In each, the network is a layer that routes value flowing through it from elsewhere. **In Cascade, the network is the economy itself.**

6.2 Augmentation, Not Replacement

Cascade is positioned as an augmentation to existing intellectual-property regimes rather than a replacement for any of them. Copyright, patents, trade-secret protections, sui generis database rights, moral-rights frameworks, and analogous statutory mechanisms continue to apply to works registered with the system, and creators retain the full set of rights and remedies available under any regime that governs their work. Registration does not constitute a waiver, grant, license, transfer, or surrender of any existing right.

The framing is structural. Existing regimes do not provide a positive economic mechanism for channeling value through the acknowledgment of influence. Cascade supplies that function. A creator may simultaneously enforce copyright, license a work for derivative use, and receive citation-driven compensation through Cascade — each mechanism operating in its own domain without interference. The system layers on top of existing law; it does not require existing law to change.

6.3 Aligned Incentives

Cascade structurally aligns patron incentives with influence propagation rather than scarcity extraction. A patron's return depends on the work's *generative capacity* — the likelihood that the work becomes an input to subsequent creation — because subsequent citations and patron purchases of citing works route value upstream through the cascade to the patron's selected target. This is the opposite of secondary-market scarcity systems, where a patron's return depends on resale price in markets external to creation itself.

Citation honesty is similarly aligned. A creator who under-cites prior works retains a larger creator allocation, but the uncited creators have an economic interest in being cited that they can pursue through community-proposal-and-voting or AI-curation review periods. Economic stakes are present on both sides of the citation question, which is what makes the citation-curation governance mechanisms in Section 4.7 enforceable in the first place.

6.4 Progressively Self-Completing

Cascade is progressively self-completing. The ghost-node mechanism extends the citation graph backward in time as off-chain creators surface to claim accumulated escrow, and each claim potentially extends the graph further back through the citations of the newly-claimed work. The network does not require comprehensive initial population to function — it back-fills as participation grows, with economic incentive carrying the back-filling rather than persuasion or platform outreach. Over time, the system's representation of human creative influence becomes more complete as a structural consequence of how participation works.

7 Topology-Driven Discovery

The reflexive property described in Section 5.3 makes the citation graph more than a routing structure. It is a substrate that systems reading it can analyze for structural features whose meaning is implied by the configuration of nodes and edges rather than by any individual node's content. This kind of analysis — *topology-driven discovery* — produces a different kind of output than the cascade does. Where the cascade propagates value through the existing graph, topology-driven discovery surfaces opportunities the graph implies but has not yet articulated.

The discovery mechanism is borrowed from work on Reflexive Knowledge Graphs, where it was originally developed for research-thesis discovery from accumulated graph topology. Applied to Cascade's citation graph, the same mechanism produces a different class of surfaced features, but the underlying operation is the same: the graph is read for structural patterns that imply something not stated in any single node.

7.1 Structural Features That Imply Meaning

Several features of the citation graph are amenable to topology-driven discovery.

Dense clusters of mutually-citing works imply that a domain has developed enough internal structure to support a synthesis or summary work that none of the individual contributors has produced. The cluster's density encodes the existence of a thesis that the works collectively support but no individual work states.

Bridges between distant regions of the graph — a single citation edge connecting two clusters that were otherwise disconnected — imply a cross-domain adjacency that has been recognized but not yet exploited. The bridge's existence indicates that a work synthesizing the two domains is possible and that someone has begun to notice.

Gaps in citation patterns — regions where the network's structure predicts works should exist but none have been registered — imply unexplored territory. The gap's position relative to surrounding density indicates the kind of work that would naturally fill it.

Centrality shifts over time indicate works whose structural importance is growing faster than their patronage reflects. A work whose centrality is rising before its share volume has caught up is structurally undervalued by the patron market — a signal both for patrons and for governance.

Each of these features can be computed from the citation graph alone, using only network-internal data. No external signal is required to detect them.

7.2 What Discovery Produces

Surfaced features can be treated as *candidates* — proposals for what the network is implying. A candidate is published in a review queue with its supporting topological evidence: the cluster, the bridge, the gap, the centrality trajectory. Community participants evaluate the candidate and decide whether to formalize it as a node in the network.

The mechanism by which a candidate gets formalized varies by implementation. In some cases, a participant accepts the candidate and authors the implied work themselves, registering it through the normal process. In other cases, the network treats the candidate itself as a kind of work — a *discovered* work registered through a different path than normal creator registration. The latter case raises a structural question about ownership that the architecture must answer.

7.3 Ownership of Discovered Nodes

A node registered through normal creator-registration has a clear creator who receives the creator allocation. A node surfaced by topology-driven discovery does not. The discovery is a function of the graph's structure and the discovery algorithm reading it. Who, then, holds the creator allocation of a discovered node?

Several distributions are plausible, and the right answer is likely community-specific.

The community as a whole. The creator allocation is allocated to the network's governance pool or distributed to participants pro-rata according to a participation metric. This treats the discovery as a collective product — surfaced by the graph that the community collectively built.

A validator subset. The creator allocation is allocated to the participants who validated the candidate during a review period — those who staked on the candidate's significance, contributed evidence supporting or refining it, or signed an attestation of its meaningfulness. This treats discovery as the joint product of the algorithm that surfaced the candidate and the humans who attested its significance. The pattern echoes the *researcher-surprise* attestation used in Reflexive Knowledge Graph research-discovery work, where human evaluation is the validating signal that distinguishes a real find from a hallucinated one.

The treasury. The creator allocation flows entirely to the treasury, where governance disburses it. This treats the discovery as a public-good output of the network and channels the value to whatever activities governance has authorized.

A hybrid split. The creator allocation divides across multiple of the above — some portion to the community, some to validators, some to the treasury — with the split itself a governable parameter.

Citation and treasury allocations of the discovered node operate as they would for any other node: cited prior works receive the citation allocation that cascades upstream, and the treasury allocation reaches the treasury as in any registration. The unsettled question concerns only the creator allocation, and only when no individual creator exists.

Different communities operating Cascade implementations will likely choose differently. A research-oriented community may favor validator attestation, where the validators are field experts whose attestation carries epistemological weight. A community building a shared public-infrastructure may favor a treasury-funded model where discovered work is treated as public infrastructure. A creative-community implementation may favor a community-pro-rata model where every participant is treated as a contributor to the structural conditions that produced the discovery.

The architecture supports all of these. What the architecture does *not* do is force a single answer. The question of who owns a discovered node is, like the question of who decides what counts as influence (Section 4.7), a design parameter rather than a fixed assumption.

8 Considerations and Open Questions

Several aspects of the architecture are deliberately left underspecified, either because they admit multiple reasonable answers that communities must choose among, or because they raise questions that would benefit from real-world deployment before being committed. This section enumerates the most significant.

8.1 Citation Quality

The architecture aligns economic incentives with honest citation, but alignment is not enforcement. A creator who under-cites still retains a larger creator allocation; the uncited creators must successfully challenge through whichever governance procedure the implementation supports. The strength of citation-quality governance — the speed and accessibility of challenge mechanisms, the burden of proof required, the cost of pursuing claims — determines how robust the citation graph is in practice. An implementation with weak governance can produce a citation graph that nominally exists but does not reflect actual influence relationships. The architecture supports robust citation-quality governance; it does not enforce it.

8.2 Calibration of Configurable Parameters

The architecture exposes several configurable parameters — the allocation function among creator, citations, and treasury; the propagation depth limit and per-hop decay factor; the treasury allocation percentage; the citation-curation mode policies; the discovery-node ownership distribution. Each parameter shapes system behavior in significant ways. The right calibration for a given community depends on the kinds of works being registered, the cadence of patronage, the depth of typical citation chains, and the community's preferences about how value should distribute. Initial parameter settings will be informed but not determined by deployment experience, and communities should expect to revise them through governance as the network matures.

8.3 Quality at Scale

As participation grows, the citation graph could be flooded with works of low merit or with citation patterns optimized for capture rather than acknowledgment. The architecture's resistance to flooding rests on citation-curation governance and on the economic friction of registration. As deployment scales, the calibration of those mechanisms — registration costs, citation-quality thresholds, AI-curation rules — will need to keep pace. The question of how a community maintains signal-to-noise in a growing graph is not solved at the architectural level; the architecture provides the levers, and governance decides where they sit.

8.4 Identity and Sybil Resistance

Several mechanisms in the architecture depend on stable participant identity — patron-share purchases, ghost-node conversions, citation-quality challenges, treasury governance. The strength of an implementation's identity layer determines its resistance to Sybil attacks, in which a single actor controls many ostensible identities to skew outcomes. Cascade does not specify an identity scheme and is compatible with any — self-sovereign-identity protocols, public-key attestation, platform-anchored identity, federated systems — but the choice has significant downstream effects on the network's robustness against gaming. This is one of the most important decisions an implementing community makes.

8.5 Implementation Neutrality and Cross-Implementation Interoperability

The architecture deliberately does not specify particular blockchain technology, governance form, identity scheme, or content-identification system. Different communities will instantiate Cascade in different ways. The mechanisms described in this whitepaper specify the architecture's behavior at the level of structural commitments; the technical substrate that supports those commitments is a choice each implementation makes. This neutrality is intentional — the architecture is meant to operate as a substrate-level pattern rather than a particular product — but it raises the practical question of how different Cascade implementations interoperate, share lineage across ledgers, or reference each other's citation graphs. Cross-implementation interoperability is an open area not addressed by the current architecture.

9 Conclusion

Cascade describes an architecture in which the act of citing a prior work — the act of acknowledging influence — is the mechanism by which economic value flows. Citation is not metadata. It is the wire through which patronage of a downstream work compensates the upstream creators whose work made it possible, recursively, across multiple generations of influence, in a single cascading transaction.

The architecture unifies provenance and compensation. The mechanism that records *who built on whom* is the same mechanism that flows value among them. No separate enforcement layer is required because the wire is internal to the network. No external valuation signal is required because value is determined by structural position. No comprehensive initial population of the citation graph is required because the network back-fills as off-chain creators surface to claim what has accumulated for them. Patrons earn returns from a work's likelihood of being cited by future work rather than from its scarcity on external markets. Creators have an economic interest in citing honestly because the upstream creators have an economic interest in being cited. The two halves of an honest acknowledgment now point in the same direction.

The architecture is positioned as an augmentation to existing intellectual-property regimes rather than a replacement for any of them. Copyright, patents, trade-secret protections, and statutory royalty schemes continue to operate. Cascade supplies a function that none of those regimes provide: a positive economic mechanism for channeling value through the acknowledgment of influence, operating in parallel with adversarial enforcement systems without preempting their protections.

Beneath the value-propagation mechanic, the citation network functions as a typed knowledge graph — a legible substrate that supports operations beyond payment routing. Topology-driven discovery surfaces structural features the graph implies but has not yet articulated. Different communities will instantiate the architecture differently, with the configurable parameters and curation modes serving as the design surface they tune to their work.

What changes if this works is not the disappearance of existing intellectual-property regimes but the addition of an economic layer that has not previously existed. The category of intellectual contribution that has no good vehicle today — work that builds on a long lineage of public, scholarly, or open-source influences and produces value that current regimes do not route back to those influences — finally has one. **Influence becomes a thing the economy actually compensates, instead of a structural liability the economy tries to police.**