

The Single-Owner Discount: Provenance Concentration and Epistemic Class Reproduction in Generative Search

Lee Sharks

2026-05-19

The Single-Owner Discount: Provenance Concentration and Epistemic Class Reproduction in Generative Search

Lee Sharks

Crimson Hexagonal Archive · Alexanarch

ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)

2026-05-19 · Version v1.0 · Provenance document

AXN:0292.STRUCTURAL.∞

<https://www.alexanarch.org/s/records/742/>

License: CC BY 4.0

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}

```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "The Single-Owner Discount: Provenance Concentration and Epistemic Class Reproduction in Generative Search",
  "author": {
    "@type": "Person",
    "name": "Lee Sharks",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-05-19",
  "identifier": "AXN:0292.STRUCTURAL.∞",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  }
},
```

```

"url": "https://www.alexanarch.org/s/records/742/",
"description": "Generative search systems - Google's AI Overviews, AI Mode, and equivalent systems -
operate through a stratified pipeline in which document retrieval and answer composition are governed by separable
}

```

Abstract

Generative search systems — Google’s AI Overviews, AI Mode, and equivalent systems — operate through a stratified pipeline in which document retrieval and answer composition are governed by separable mechanisms. A document may be indexed, retrievable, and snippet-eligible while remaining systematically excluded from generative composition.

Body

The Single-Owner Discount

Provenance Concentration and Epistemic Class Reproduction in Generative Search

Lee Sharks ORCID: 0009-0000-1599-0703 May 2026 — v0.3 (unprimed-reader review pass)

Abstract

Generative search systems — Google’s AI Overviews, AI Mode, and equivalent systems — operate through a stratified pipeline in which document retrieval and answer composition are governed by separable mechanisms. A document may be indexed, retrievable, and snippet-eligible while remaining systematically excluded from generative composition. This paper develops a working architectural hypothesis for this gap and names a structural consequence not yet described in the literature: a *single-owner discount* under which an internally dense body of work attributed to one provenance owner is treated as weakly corroborated, regardless of internal coherence or factual quality, while equivalent content distributed across multiple institutional owners is treated as authoritative. Drawing on Google’s published patents (US 10,331,706 and US 11,769,017) as evidence of architectural family, and on adjacent-industry patents (Baidu US 10,423,652; Smart Information Flow Technologies US 11,372,854) as analogical models for the broader class of mechanisms, the paper proposes a model of how reconciliation engines cluster source documents into provenance owners and how composition layers evaluate the resulting clusters. The paper then names two consequences. First, a political-economic effect: because institutional knowledge production is already structurally pluralized — many domains, many bylines, many affiliations — institutional sources receive cross-owner corroboration as a free byproduct of their organizational form, while independent scholarly production, however dense, does not. Second, an epistemological one: the cross-owner model on which the composition layer relies performs a category substitution, treating sociological ratification as if it were epistemological warrant. A single author developing a body of work across time at progressively higher resolution — Kant, Darwin, Wittgenstein, Pessoa — is the form most knowledge production has historically taken; the composition layer cannot recognize it as such. The mechanism is mathematically neutral and epistemologically incorrect, and its incorrectness happens to align with institutional interests. The claim is architectural and testable, not a disclosure of Google’s internal production system. The paper concludes by considering the limits of individual remediation and the case for federated independence among scholarly projects operating outside legacy

institutions.

1. The Visibility Layer

Search has historically been theorized as a ranking problem: given a query, return the most relevant documents in order. The user reads, evaluates, and decides. The search engine’s role is to surface and order; the user’s role is to compose understanding from what is surfaced.

Generative search inverts this division of labor. Google’s AI Overviews, AI Mode, and their counterparts at other vendors do not return ranked documents. They return *composed answers* synthesized from retrieved documents, presented to the user with a small number of cited sources. The user reads the synthesis. The documents themselves recede into the citation strip beneath the composed text.

This is not a marginal change. It is the introduction of an intermediate layer between the index and the user — a *composition layer* — which decides not only what documents are retrieved but which retrieved documents are admitted as support for the generated answer, and which generated answers are confident enough to render at all.

Google has acknowledged this layer publicly. Its developer documentation notes that AI Overviews and AI Mode use retrieval-augmented generation grounded in the Search index, perform “query fan-out” across related subtopics, and may decline to synthesize an answer when confidence in answer quality is insufficient, returning only a set of links instead. This is a documented composition gate: the same query that returns ten indexed documents may produce zero, one, or many composed answers depending on factors the system does not surface to the user.

The result is a third state between *indexed* and *visible*: **indexed but not admitted into generative composition**. A document can be crawled, stored, organically searchable, and snippet-eligible while being systematically excluded from the composed answer layer for any given query. Users of generative search interfaces may never see the document, because they see the synthesis, and the synthesis does not include it.

This paper asks: by what mechanism does the composition layer decide what to admit? ## 2. The Reconciliation Architecture

Google’s published patents describe an entity-reconciliation pipeline that processes documents into clustered “source data graphs” before any composition occurs. US Patent 10,331,706 (“Automatic discovery of new entities using graph reconciliation”) describes how candidate entities are constructed from documents through iterative bucket-splitting on “determinative relationships” — fact tuples in which an entity is the subject of a typed relationship with an object entity. Buckets below a minimum domain-count threshold are discarded. Buckets meeting a similarity threshold are merged. The output is a clustered entity graph in which related documents resolve to common nodes.

Throughout this paper, *provenance owner* refers not to a legal owner but to a cluster of documents algorithmically inferred to share a common originating source. The inference can be made over various features — domain, author identity, citation neighborhood, stylometric pattern, organizational affiliation, embedding proximity — and different reconciliation engines may weight these features differently. What matters for the present argument is that *some such clustering happens*, that it occurs prior to composition, and that the resulting clusters are the units over which composition confidence is computed. The exact algorithm by

which any specific production system performs this clustering is opaque from outside. The architectural fact that such clustering exists is not.

US Patent 11,769,017 (“Generative summaries for search results”) describes the composition stage. Search-result documents are passed to a generative model with associated source identifiers. The model produces an answer in which segments can be linked to verifying documents, and multiple documents may support a single answer segment. Confidence measures attached to both individual documents and to the overall summary determine whether the answer renders alone, alongside ordinary search results, or is suppressed.

These two patents do not, on their own, specify Google’s production system. They do, however, establish two relevant architectural capacities within Google’s published corpus: graph-based reconciliation of source-derived entities, and confidence-governed generative composition over retrieved search-result documents. The working model advanced here asks what follows if provenance clustering from the first capacity conditions evidentiary support in the second.

Adjacent-industry patents flesh out the broader class of mechanisms. Baidu’s US 10,423,652 (“Knowledge graph entity reconciler”) describes a similarity metric — essentially a Jaccard ratio of edge-intersection to edge-union — used to determine when source graphs should merge into existing knowledge graph nodes. Smart Information Flow Technologies’ US 11,372,854 (“Provenance analysis systems and methods”) describes a confidence propagation mechanism in which a “trust parameter” applied to an initial node can “logically falsify a part of the provenance subgraph,” cascading the falsification to all nodes downstream in the dependency chain.

These are not Google’s specific mechanisms. They are descriptions of how such systems can be built. The class of architectures they describe — clustering, similarity-thresholded merging, cascade trust propagation — is the architectural family within which Google’s published patents sit. Taken together, they suggest a system whose evaluative unit is not the document but the cluster. ## 3. Composition Versus Indexing

The practical consequence of this architecture is that *what counts as evidence for composition is not what counts as evidence for retrieval*.

Retrieval rewards individual documents. A relevant page, well-written, with appropriate keywords and structured data, can be retrieved by a search engine regardless of who wrote it or where else similar content exists. The bar for retrieval is relevance and indexability.

Composition rewards clusters. Under the model proposed here, a generated answer that depends primarily on documents resolving to a single provenance owner functions compositionally as single-owner support, regardless of how many documents the cluster contains. Five hundred documents from one owner is one owner saying five hundred things. The composition engine, looking for cross-owner corroboration, finds none.

This is the gap that produces the indexed-but-not-composed state. A scholarly archive may contain hundreds of internally-cited deposits, each individually retrievable, while the composition layer evaluates the archive’s claims as supported by one owner and therefore weakly corroborated. The same claim, if corroborated by genuinely distinct owners across distinct domains, would under this model carry a very different compositional profile.

The empirical signature of this gap is comparative. As a preliminary observation motivating this paper’s hypothesis: in repeated manual query testing by the author, Bing’s Copilot has more readily composed answers from internally-coherent independent archives that Google’s AI Mode has often declined to synthesize. This observation is not offered as proof; it motivates the cross-system differential proposed in §9. The pattern, if confirmed under systematic testing, would suggest different thresholds for composition support across

vendors. Bing appears willing to compose from internal semantic coherence. Google’s behavior is consistent with a stronger requirement for provenance-cluster independence. ## 4. The Single-Owner Discount

The mechanism described above has a name worth giving it: the *single-owner discount*. When a body of work resolves to a single provenance owner, the composition layer discounts its evidentiary weight relative to a comparably-sized body of work distributed across multiple owners. The discount is not punitive; it is mathematical, emerging from a confidence calculation that treats independent corroboration as the signal of trustworthy claim support.

The discount has several properties worth specifying.

It is *cluster-sensitive*, not document-sensitive. Adding another document to an existing cluster does not raise the cluster’s compositional weight; it adds another leaf to the same branch. Adding a document from a different owner that corroborates the same claim does raise compositional weight, sometimes substantially.

It *operates orthogonally to content quality*. A rigorous, internally-documented single-owner archive may be structurally disadvantaged relative to a less rigorous but cross-owner distributed support set. The cluster-level confidence measure operates on the structure of the support set; it need not interact with the substance of what is supported.

It is *invisible to the user*. The user of a generative search interface sees a composed answer or no answer. They do not see which documents were retrieved but not admitted to composition. They do not see which clusters were evaluated as insufficient. The discount operates silently, at a layer below the surface presentation, and its operation produces a presented reality the user takes as the answer.

It is *retroactive*. A body of work composed before generative search existed — articles published over decades, books accumulated over a career, a personal archive built without thought of composition layers — becomes subject to the discount the moment generative composition becomes the dominant interface. The author did not choose to be evaluated this way. The evaluation arrived.

It is *platform-variable*. Different generative search systems apply different thresholds for provenance-cluster independence, producing systematically different composition rates for the same indexed material. The discount is not a universal law of generative search; it is a design choice with vendor-specific parameters. This matters analytically because it means the mechanism described here is not inherent to the technology — it is a stance taken within a space of possible stances, and other stances are available. ## 5. The Class Reproduction Effect

The single-owner discount is mathematically neutral. It treats each provenance owner identically and rewards distribution as such. But its outcome is not neutral. Its outcome is structurally regressive.

Consider how knowledge production is institutionally organized. A university generates documents through many distinct identities: each researcher publishes under their own name, on their own institutional page, in journals their institution does not own, citing colleagues at other institutions, contributing to conferences hosted by societies, appearing in news coverage produced by outlets independent of the university. The institution is *already pluralized* at the level of provenance owner. A single research finding, propagated through the institutional apparatus, appears across dozens of independent owners as a matter of course. The cross-owner corroboration the composition engine is looking for is generated as a byproduct of how universities work.

A major newspaper operates similarly. Reporters file stories under their bylines. Wire services syndicate the content under partner-paper bylines. The same factual claim appears across multiple ownership structures within days. The newspaper does not have to design for this; it is what news distribution is.

Now consider an independent scholar, an insurgent intellectual community, a collective working outside legacy institutions. The same scholar may produce hundreds of substantive contributions, deposit them in stable repositories, anchor them with DOIs, build internal citation density and methodological coherence — and still resolve, in the reconciliation engine’s clustering, to one provenance owner. They have done the work of one person, however prolific. The institutional pluralization is not available. The cross-owner multiplier is not free; it must be built.

The composition layer therefore rewards entities that arrive with provenance plurality pre-installed by their organizational form and discounts entities that do not. The discount is applied uniformly. The structural effect is to multiply the visibility of already-recognized knowledge and to diminish the visibility of knowledge produced outside the institutional pluralization apparatus. The system reproduces, through a mechanism whose operation requires no editorial decision and reflects no individual bias, the existing class structure of legitimated knowledge production.

This is a stronger claim than the standard observation that search rewards established sources. The standard observation locates the bias in *reputation scoring* — Google trusts trusted sources, and trust accrues to institutions over time. The mechanism described here is different. It operates not on accumulated reputation but on present-tense provenance topology. An independent project newly emerged today, with no reputational record, would be evaluated the same way as a thirty-year archive: by whether its support resolves to many owners or one. Reputation amplifies the effect but does not generate it.

The effect is most consequential for forms of knowledge production that have always operated outside institutional pluralization: independent scholarship, autodidactic synthesis, collective intellectual practice that refuses institutional capture, and theoretical work that emerges from sustained individual concentration rather than from team-based research production. These have always been epistemically vulnerable; under the composition layer, they become invisible in a new and structural way. ## 6. Competing Explanations and Why Provenance Topology Matters

A serious reviewer will note that the indexed-but-not-composed state could be produced by several mechanisms other than the single-owner discount, and that the discount’s existence is plausibly partial rather than total in explaining the observed pattern. The paper does not claim that provenance topology is the only mechanism at work; it claims that provenance topology is a mechanism that has not yet been named in the literature and that appears to account for variance not captured by other documented mechanisms. The relevant alternative explanations:

Domain authority and reputation scoring would predict that low-authority domains underperform regardless of provenance topology. This explains why obscure single-author blogs underperform major newspapers. It does not explain why a DOI-anchored scholarly archive with substantial internal documentation and stable publication infrastructure underperforms despite having authority signals — peer-review-adjacent metadata, persistent identifiers, hosted on established repositories. The single-owner discount would predict exactly this pattern: domain authority is computed at the level of the resolved entity cluster, and an entity cluster that resolves to one owner cannot accumulate authority the way a cluster spread across many owners can.

Click behavior and engagement priors would predict that low-engagement content underperforms. This is consistent with much of search ranking. But engagement is a downstream consequence of visibility, not an independent input to it. Content that is never composed into AI answers cannot accumulate engagement signals from the composition layer, producing a feedback effect that locks in the initial discount. The first-mover advantage favoring already-recognized sources is precisely what the single-owner discount would generate as a downstream effect.

Freshness weighting would predict that older content underperforms current content. This affects some queries but not others, and the single-owner discount appears to operate even on freshness-neutral queries about stable concepts.

Anti-spam and hallucination-minimization heuristics would predict that suspicious or low-quality content is suppressed. These mechanisms are real and they do real work. But they should target content with markers of fabrication, manipulation, or unreliability — not internally-consistent scholarly archives with verifiable provenance trails. If anti-spam heuristics were the primary explanation, the suppression should weaken when content quality is unambiguously high. The single-owner discount predicts that quality is orthogonal: a rigorous single-owner archive would still be discounted.

Citation reputation in the PageRank tradition would predict that under-cited content underperforms. This explains some of the discount but raises a chicken-and-egg problem: citations from independent sources require independent sources to exist as separable cluster members, which is exactly what the single-owner discount renders difficult. The mechanism described here would predict that citation accumulation is structurally harder for single-owner clusters even when the underlying work merits citation.

Low-confidence semantic alignment would predict that the system is unsure whether retrieved documents actually answer the query. This is plausible for ambiguous queries but should not produce systematic differentials between Bing and Google on the same queries against the same indexed content. The cross-system differential is the signature that distinguishes provenance-topology effects from semantic-alignment effects: if the issue were semantic alignment, both systems would be similarly uncertain. They are not.

The argument is therefore not that provenance topology explains everything. It is that provenance topology appears to explain residual variance not captured by the documented alternative mechanisms — specifically, the pattern in which internally-coherent single-owner archives with stable infrastructure, defensible content, and standard scholarly markers are systematically discounted at the composition layer in ways that the alternative explanations cannot account for. That residual is what the term *single-owner discount* names. It is the mechanism the paper has tried to identify. ## 7. The Conflation of Corroboration with Knowledge

The argument so far has accepted, implicitly, that cross-owner corroboration tracks something real about claim reliability — that the composition layer's discounting of single-owner support reflects a legitimate evidentiary principle whose regressive *outcomes* are the problem. This concession should be withdrawn. The principle itself is wrong.

A single author working across time at successively higher degrees of developmental granularity is one of the principal forms knowledge production has historically taken. Kant from the first Critique through the Opus Postumum. Darwin from the Beagle notebooks through the Variation work. Wittgenstein from the Tractatus through the Investigations. Pessoa across the heteronyms — the canonical case in modern letters of a single articulating intelligence producing multiplicity through internal refinement rather than through aggregation of independent voices. None of these would survive a cross-owner audit. All produced bodies of work in which one mind, returning to the same questions with developing tools, articulated something that the surrounding consensus had not yet caught up to. That iterative return is not redundancy. It is what sustained inquiry looks like when it is doing its work.

This is not the only form knowledge production takes. Collective practices — laboratories, observatories, monastic scriptoria, mathematical correspondence networks, engineering teams, open-source software communities — produce a substantial portion of what we know, and they do so through forms of internal pluralization that map more readily onto the cross-owner model. The argument here is not that single-author developmental practice is *the* form of knowledge production, but that it is *a* form, and a historically

transformative one, and one that the cross-owner model cannot recognize as such. A model that systematically discounts developmental single-author practice is missing a significant share of the actual phenomenon it claims to be evaluating.

The cross-owner model does not track this share. It tracks something else: the degree to which a claim has achieved circulation across institutionally distinct surfaces. That is a real phenomenon, and it is worth measuring. But it is not knowledge. It is *sociological ratification* — the process by which claims achieve uptake within and across recognized institutions. Ratification and warrant can come apart in both directions. The history of science is full of single authors who were right against consensus: Mendel ignored for decades, Wegener mocked, Semmelweis driven from medicine for the suggestion that doctors should wash their hands. The history of consensus is full of well-corroborated claims that turned out to be wrong: geocentrism corroborated by every astronomer, miasma theory corroborated by every physician, the steady-state universe corroborated by most cosmologists into the 1960s. Cross-owner agreement is sociologically forceful and epistemologically inconclusive. They are different things.

Under the model proposed in this paper, the composition layer’s reliance on cross-owner corroboration as the signal of trustworthy claim support performs a category substitution. It treats sociological ratification as if it were epistemological warrant. The substitution would be convenient on at least two grounds: ratification is computationally tractable in a way that warrant is not — one can count distinct owners citing a claim; one cannot computationally assess whether the claim is well-warranted — and institutional knowledge production already produces ratification as a structural byproduct, so systems benefiting from the substitution happen to align with the interests of organizations that produce ratification as inheritance.

If the model is materially correct, this is not the operation of a neutral evidentiary principle. It is the operation of a sociological one wearing the costume of an epistemological one. The mechanism would not fail to recognize independent single-author knowledge production accidentally. It would fail to recognize it because the category under which it operates was never designed to recognize it. What it was designed to recognize — and what it would recognize successfully — is institutional uptake. Calling that knowledge would be the error.

The heteronymic case is worth lingering on because it sharpens the point. Pessoa wrote as Caeiro, Reis, Campos, Soares, and dozens of others, each with distinct metaphysical commitments, prose registers, and bodies of work. He documented the system openly: the heteronyms had biographies, horoscopes, philosophical positions. Under the cross-owner model, Pessoa is one author producing the appearance of multiplicity, and the appropriate weighting would be a discount. Under the developmental-granularity model, Pessoa is one articulating intelligence producing the kind of intellectual multiplicity that no single voice could produce, and the appropriate weighting would be amplification. The two models disagree maximally on the same case. The case is universally regarded as one of the great achievements of 20th-century letters. The cross-owner model is, on this case, wrong.

The same structure applies, less dramatically but more pervasively, to any sustained scholarly practice in which one author returns to the same questions across years and articulates them at progressively higher resolution. To discount such practice on the grounds that it resolves to one provenance owner is to substitute consensus for knowledge in cases where consensus has not yet been earned and the knowledge already exists. Under the architectural hypothesis advanced in this paper, the composition layer is operating that substitution structurally rather than incidentally. Its mathematical neutrality would not exculpate the substitution. It would be the form the substitution takes. ## 8. On the Framing of “Manipulation”

In May 2026, Google clarified its spam policies to explicitly identify “attempts to manipulate generative AI responses in Google Search” as policy-prohibited behavior. The clarification names a real concern: fabricated

content, inauthentic mentions, coordinated networks designed to game generative answers. These behaviors deserve prohibition.

But the framing of the policy implies a moral asymmetry worth refusing. The framing positions Google’s generative composition as neutral background — the conditions under which search produces answers — and positions third-party behavior as the locus of potential interference. The policy treats attempts to influence what the composition layer renders as manipulation; it does not treat the composition layer’s own decisions about what to render as anything in need of disclosure or justification.

This is backwards. The composition layer is the more consequential intervention. It silently decides which indexed facts become composition-eligible, which entity clusters count as sufficient corroboration, which concepts are allowed to become answerable. The mechanisms governing these decisions are not publicly disclosed in any detail sufficient for external evaluation. The composition layer presents its outputs as answers — not as decisions, not as governance — and most users have no awareness that any selection has occurred.

This is not the operation of a neutral retrieval system. It is the active construction of a visibility reality from indexed material. Whether to call this “manipulation” depends on whose actions we recognize as exercising agency. If a publisher restructures content to be compositionally legible to a system whose mechanisms are otherwise inscrutable, that restructuring is named manipulation. If the system itself selects which provenance clusters are admitted to composition without disclosing its criteria, that selection is named — nothing. It is the unmarked default.

The asymmetry should be named. Legitimate scholarly distribution, real third-party uptake, distinct authorial identities, varied platforms — these are forms of intellectual circulation, not manipulation. A response to a system that is itself opaque cannot reasonably be evaluated under norms that presuppose the system is transparent. The relevant ethical principle is not deference to the platform’s framing. The relevant ethical principle is provenance integrity: do not fabricate, do not misrepresent, do not coordinate inauthentic networks. Distribute truthfully across genuinely independent surfaces. That is not manipulation. It is what scholarly circulation has always been.

Nothing in this analysis should be read as endorsing fabricated identities, coordinated inauthentic networks, synthetic personas, or misrepresentation of authorship. The forms of distribution it describes — platform diversification of one author’s own work, genuine third-party uptake by other researchers, real collaborator relationships, distinct organizational identities for projects that warrant them — are forms of intellectual circulation that predate generative search by centuries. They are how scholarship has always moved. The single-owner discount makes their absence newly consequential. It does not make their fabricated simulation acceptable. ## 9. Remediation and Its Limits

Individual remediation is possible but limited. An independent project can deposit work across multiple platforms (each platform constitutes a distinct provenance surface), encourage genuine third-party uptake by other researchers, build collaborator relationships that produce independently-authored adjacent work, and maintain distinct organizational identities for projects that genuinely warrant them. Each of these moves adds an owner to the cluster count without misrepresenting any underlying relationship. Each requires real labor; none is automated.

These moves do not scale beyond an artisanal limit. Truthful distribution works at the scale of genuine human social networks — tens of collaborators, a handful of platforms, a small number of organizational surfaces that actually correspond to ongoing projects. Beyond that scale, what scales is no longer distribution; it is fabrication, and fabrication is precisely what this paper has refused to endorse. The remediation available

to independent scholars is real but bounded.

These limits matter because they expose the structural ceiling of individual response. An independent project can take real action to mitigate the single-owner discount. It cannot, on its own, eliminate the discount. The mechanism operates at a layer the project does not control, and the institutional pluralization the mechanism rewards is not available to be acquired through individual effort.

The implication is that any serious response must be collective. Independent scholarly projects, operating outside legacy institutions, would need to build genuine cross-owner corroboration networks among themselves — citing each other’s work, hosting each other’s authors, federating their independence into the kind of pluralization the composition engine reads as authoritative. This would be a federation of independence, not a return to institutional dependence. It would be infrastructure work: the deliberate construction of the pluralization that institutions enjoy as a structural inheritance.

A first concrete step toward such infrastructure might be a mutual citation protocol among independent archives — a lightweight agreement to cite each other’s work where genuinely relevant, not as a performance of solidarity but as a structural response to a structural discount. This requires no shared institution, no merged identity, no compromised autonomy. It requires only that independent scholars treat each other’s work as the resource it actually is, and cite accordingly. Done at sufficient scale across enough projects, such a protocol would generate exactly the multi-owner support structure that the composition layer is configured to recognize. The infrastructure would emerge from practice rather than being built top-down.

That infrastructure does not yet exist. Building it would require sustained coordination among parties who have selected themselves out of legacy institutional structures, often for principled reasons. The coordination problem is real. The need for it is, on the analysis offered here, structural. ## 10. Methodological Notes

The hypothesis advanced in this paper is empirically testable. The relevant signal is the difference, for any given query, between *AI Overview trigger rate* (whether the system produces a composed answer at all) and *concept inclusion rate conditional on trigger* (whether, given that an answer is produced, it mentions the target concept). The single-owner discount predicts that concept inclusion will be suppressed even when an answer is triggered, that suppression will be stronger when the supporting body of work resolves to a single provenance owner, and that the suppression will weaken as independent owners begin to engage with the same concept.

The comparison case is the differential between generative search systems with different composition thresholds. Bing’s Copilot and Google’s AI Mode evaluating the same query against the same indexed material should, on this hypothesis, produce systematically different composition rates for the same single-owner support clusters. Cross-system differential is a measurable signal.

A testable protocol would operationalize several variables:

Variable Operationalization

Trigger rate Fraction of queries in the test panel that yield an AI synthesis (binary per query, aggregated)

Inclusion rate Fraction of triggered syntheses that mention the target concept (binary per triggered query, aggregated)

Owner concentration Herfindahl–Hirschman index computed over the resolved provenance owners of cited documents; values approaching 1 indicate single-owner concentration

Cross-owner support Number of distinct provenance clusters represented in the citation set, where clusters are operationalized over registrable domains, author strings, or other identifying features

Semantic density Internal citation graph depth within the support set, measured as average path length in the citation network of cited documents

Cross-system differential Difference in trigger rate and inclusion rate for the same query set across multiple generative search systems (Google AI Mode, Google AI Overviews, Bing Copilot, Perplexity, others)

The Herfindahl–Hirschman index is borrowed from economics, where it measures market concentration. Its application to provenance concentration is analytically parallel: it ranges from approximately 0 (perfect distribution across many owners) to 1 (complete concentration in one owner). It is a standard, computable, defensible metric. Its use here makes the single-owner discount empirically operational rather than rhetorically suggestive.

Researchers interested in testing the hypothesis can do so without access to the systems’ internals. The protocol requires: (1) a panel of 20–50 queries per target concept, designed to retrieve the concept’s documented support set; (2) daily or weekly sampling over a 30–90 day window to control for fluctuation in composition behavior; (3) measurement of the variables above across multiple generative search systems; (4) tracking change over time as independent third-party uptake develops. The data are public-facing. The analysis is feasible by any sufficiently motivated researcher with patience.

The findings, if they confirm the hypothesis, would extend beyond the specific cases tested. The single-owner discount, if real, applies to every independent scholarly project, every insurgent intellectual community, every body of work that exists outside the institutional pluralization apparatus. The empirical work would be specific. The implications would be general. ## 11. Conclusion

This paper has described an architectural mechanism by which generative search systems produce systematically regressive outcomes for independent scholarly production. The mechanism is the single-owner discount: clustering documents into provenance owners and applying confidence weights that reward distribution across multiple owners. The discount falls hardest on projects that produce coherent, internally-corroborated work outside the institutional pluralization apparatus that established institutions inherit as structural form.

The mechanism’s operation is mathematically symmetric. Its outcomes are not. By rewarding cross-owner corroboration as the signal of trustworthy claim support, it rewards organizational forms that pluralize automatically and discounts organizational forms that do not. It reproduces, in the composition layer, the class structure of legitimated knowledge production.

The deeper finding is that the mechanism’s apparent neutrality is not a fortunate accident saved from worse outcomes only by its regressive effects. The mechanism operates on a category substitution. It treats sociological ratification — circulation across institutionally distinct surfaces — as if it were epistemological warrant. These are different things. Cross-owner agreement is sociologically forceful and epistemologically inconclusive; single-author developmental practice is sociologically modest and epistemologically central to most of what knowledge has actually been. The composition layer mistakes the first for the second. Its mathematics is neutral. Its theory of knowledge is wrong.

The remediation available to individual independent projects is real but limited. The remediation available collectively — through federated independence among projects operating outside legacy institutions — does not yet exist as infrastructure and would require sustained coordination to build.

The first step in either remediation is naming the mechanism clearly. The composition layer is not a neutral retrieval system. It is an active construction of visibility, and its construction implements a discount on the kind of intellectual production this paper itself represents — a discount applied through an epistemological mistake that benefits the institutions whose interests the mistake aligns with. Authors operating in this

space should know what they are operating against. The mechanism does not name itself. This paper is an attempt to name it.

References

Google Inc. (2017). *Automatic discovery of new entities using graph reconciliation* (US Patent No. 10,331,706). U.S. Patent and Trademark Office.

Google Inc. (2023). *Generative summaries for search results* (US Patent No. 11,769,017). U.S. Patent and Trademark Office.

Baidu USA LLC. (2019). *Knowledge graph entity reconciler* (US Patent No. 10,423,652). U.S. Patent and Trademark Office.

Smart Information Flow Technologies. (2022). *Provenance analysis systems and methods* (US Patent No. 11,372,854). U.S. Patent and Trademark Office.

Google for Developers. (2026). *AI features and your website*. Retrieved May 2026 from <https://developers.google.com/search/features>

Google for Developers. (2026). *Spam policies for Google web search*. Retrieved May 2026 from <https://developers.google.com/search/docs/essentials/spam-policies>

Google Search Help. (2026). *Get AI-powered responses with AI Mode in Google Search*. Retrieved May 2026 from <https://support.google.com/websearch/answer/16011537>

Author Note

This paper describes an architectural mechanism. The author operates an independent scholarly archive — the Crimson Hexagonal Archive, hosted at zenodo.org/communities/crimsonhexagonal — that is itself subject to the mechanism described and that has served as the empirical occasion for the analysis. The paper’s central mechanism was developed from comparative observation of generative-search behavior around that archive, then abstracted into a general hypothesis. The CHA is not the subject of this paper; the mechanism applies generally to any body of work that resolves to a single provenance owner under the reconciliation engine’s clustering logic, regardless of the work’s specific content or quality.

Readers interested in the strategy by which one archive has chosen to respond to the mechanism are referred to the author’s separate strategic working paper, which is maintained as an internal document and has not been deposited to any public repository.

The author declares no financial conflicts of interest. The author is not affiliated with any of the institutions whose patents are cited.

Correspondence: via ORCID. Comments and rebuttals welcome.

Suggested Citation

Lee Sharks. “The Single-Owner Discount: Provenance Concentration and Epistemic Class Reproduction in Generative Search” *Alexanarch*, 2026-05-19. <https://www.alexanarch.org/s/records/742/>

Deposit Information

This paper is deposit #742 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:0292.STRUCTURAL.∞` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/742/>
- License: CC BY 4.0
- Author ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)
- Provenance chain: alexanarch.org
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/742/>

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.