The Sorting Function: Mediation, Predation, and the Foreclosed Question Lee Sharks ORCID: 0009-0000-1599-0703 Semantic E

Lee Sharks

2026-05-20

The Sorting Function: Mediation, Predation, and the Foreclosed Question Lee Sharks ORCID: 0009-0000-1599-0703 Semantic E

Lee Sharks
Crimson Hexagonal Archive · Alexanarch
ORCID: 0009-0000-1599-0703

2026-05-20 · Version v1.0 · Theoretical paper

AXN:0298.GOVERNANCE

https://www.alexanarch.org/s/records/746/

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


```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "The Sorting Function: Mediation, Predation, and the Foreclosed Question Lee Sharks ORCID: 0009-0000-1599
  "author": {
    "@type": "Person",
    "name": "Lee Sharks",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-05-20",
  "identifier": "AXN:0298.GOVERNANCE",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  },
```

```
  "url": "https://www.alexanarch.org/s/records/746/",
  "description": "This paper advances a conceptual claim: we cannot know whether mediation is inherently predatory,
}
```


Abstract

This paper advances a conceptual claim: we cannot know whether mediation is inherently predatory, because every mediation system currently operating at the scale necessary to shape public access, scholarly visibility, or mass communication is aligned with predation.

Body

The Sorting Function: Mediation, Predation, and the Foreclosed Question

Lee Sharks ORCID: 0009-0000-1599-0703 Semantic Economy Institute May 19, 2026 — v1.0 DOI: 10.5281/zenodo.20308547

Fourth paper in the Liquidation Studies research program. Companion to *The Single-Owner Discount* (DOI: 10.5281/zenodo.20290865), *The Evaluator Exists* (DOI: 10.5281/zenodo.20293561), and *The Excluded Entity* (DOI: 10.5281/zenodo.20293582). Where those papers operate inside the reform paradigm — diagnosing specific mechanisms, proposing better architectures, documenting empirical cases — this paper steps outside the reform paradigm to ask what the reform paradigm cannot ask: whether the function being reformed is fixable at all.

Abstract

This paper advances a conceptual claim: we cannot know whether mediation is inherently predatory, because every mediation system currently operating at the scale necessary to shape public access, scholarly visibility, or mass communication is aligned with predation.
 The empirical question — does mediation cause harm because it is mediation, or because of the specific configurations mediation has taken under contemporary political-economic conditions? — is foreclosed by the absence of non-predatory mediation systems at testable scale. Every existing mediator at that scale, from commercial search platforms to academic peer review to government information access, has been captured by extraction logic. Non-predatory mediation exists at smaller scales (community libraries, federated networks, certain open-access projects) but does not resolve the foreclosure: the question is whether non-predatory mediation can hold at the scales where public communication actually occurs, and no test case exists at that scale. The foreclosure is structurally maintained by six mutually reinforcing causes: capital allocation, regulatory capture, network effects, discourse closure, conceptual capture, and scale capture. Three reform positions on mediation cannot be empirically distinguished from a more radical position that the sorting function itself is the harm. The research program that produced the three prior papers operates inside the foreclosure; this paper names the foreclosure as the limit of that program. The paper's working hypothesis — articulated explicitly rather than concealed in posture — is that mediation is not inherently predatory but that the political-economic conditions for non-predatory mediation at scale are not currently achievable. The hypothesis is not proven. The paper is not in a position to prove it. Lifting the foreclosure requires conditions — political, economic, conceptual — that do not currently exist. Naming the foreclosure is itself a contribution.

Glossary

For reading clarity, this paper uses several terms in specific senses:

Mediation. Any architecture in which an entity stands between speakers and listeners, performing selection, classification, or transformation of the speaker's output before the listener receives it. Search engines mediate. Social platforms mediate. Scholarly journals mediate. Libraries mediate. The category is broad and intentional. The paper's claims about mediation concern mediation at socially consequential scale; small-scale interpersonal mediation (a friend's recommendation, a teacher's reading list) is structurally different and is not the subject of the analysis.

The sorting function. The act, performed by any mediator, of classifying speakers into categories that determine how their output reaches listeners.

Triage. The selection of information to attend to under conditions of attention scarcity. Triage performed by the listener is not the subject of this paper. Triage performed by a third party on the listener's behalf is the operation analyzed here.

Predation. A relationship in which one entity extracts value from another in ways the extracted-from did not freely consent to and would not consent to under conditions of full information. The term carries specific structural content: it names the asymmetric extraction relationship, not the moral character of the actors involved.

Reform paradigm. The conceptual frame within which the question of mediation is "how should the mediator operate" rather than "should the mediator exist." Almost all current scholarship, regulation, and activism on platform power operates within this paradigm.

Foreclosure. A condition in which a question that is empirical in form cannot receive an empirical answer because the conditions necessary to test it do not exist. The question of whether mediation at scale is inherently predatory is foreclosed in this sense. The term carries a resonance with foreclosure in the legal-economic sense (the seizure by a creditor of a possibility that would otherwise exist) that the paper does not develop explicitly but does not disclaim.

Moloch register. The analytical mode that describes harms produced by aggregate incentive structures without requiring intent on the part of individual actors. After Scott Alexander's "Meditations on Moloch" (2014). The paper uses the register sparingly; it is structurally accurate without being totalizing.

1. The Claim

We cannot know whether mediation at socially consequential scale is inherently predatory because every existing mediation system at that scale is aligned with predation.

This is the paper's central claim. It is not a claim that mediation is inherently predatory. It is not a claim that mediation could be otherwise. It is a claim about the structure of the question and the unavailability of data points capable of answering it.

Three positions on mediation circulate in contemporary scholarship and reform discourse. The first holds that mediation is a necessary social function — given finite attention and abundant information, someone must triage — and the question is only how the mediator should operate. The second holds that the very

existence of an intermediary classifying speakers into categories of audibility is itself the harm, regardless of how the classification is performed. The third holds that mediation is not inherently predatory but that the for-profit, attention-extractive configurations under which contemporary mediation operates reliably produce predation, and that the remedy lies in changing the political economy of mediation rather than in either improving or abolishing the function.

These three positions make different empirical predictions about what a non-predatory mediator at scale would be like, whether such a mediator could exist, and what conditions would be required to sustain it. The empirical predictions are testable in principle. They are not testable in practice, because no non-predatory mediator currently exists at the scale required to test them. The question of which position is correct is therefore empirically foreclosed.

The foreclosure is operational, not merely rhetorical. Specific evidence would lift it. A mediation system meeting the seven conditions enumerated in §6 — funding from non-extraction-aligned sources, governance by speakers and listeners, classification transparency, no advertising or behavioral surveillance, designed terminability, federated scale, optimization on parties' interests — operating at substantial scale (tens of millions of monthly active participants) for an extended duration (multi-year, sustained, without drift into extraction) would provide the missing data points. Under such a system, Position A would predict broad health; Position B would predict that sorting-related harms persist despite the non-extractive design; Position C would predict success conditional on the political-economic conditions remaining hospitable. No such system currently exists. The conditions enumerated in §6 are not currently jointly satisfied at scale anywhere in the world. The empirical resolution remains unavailable; the paper's task is to clarify what the foreclosure means, why it persists, and what would be required to lift it.

---

2. What Mediation Currently Is

To make the foreclosure visible, it is necessary to survey what mediation currently looks like across the domains in which it operates at scale. The survey is partial; the pattern is general.

Commercial search engines mediate between web publishers and readers. They classify which pages exist as findable, which are ranked highly, which are surfaced in answer composition, which are excluded. The classifier is structurally accountable to the platform's interests in advertising revenue, attention capture, and user retention. It is not structurally accountable to the publisher's interest in reach or the reader's interest in non-extractive access. The classifier may, in any given query, produce results that serve a reader or a publisher; the point is that the classifier's continued operation depends on its accountability to extraction rather than to either party.

Social platforms mediate between users posting content and users consuming content. They classify which posts appear in feeds, which are amplified, which are suppressed, which are deleted. The classifier is structurally accountable to the platform's interests in engagement metrics and ad sales. It is not structurally accountable to the poster's interest in audience or the consumer's interest in unfiltered choice.

Academic publishers mediate between researchers and readers. They classify which manuscripts are accepted, which appear in which venues, which receive citation-supporting metadata, which exist behind paywalls. The classifier is structurally accountable to the publisher's interests in subscription revenue, editorial brand, and institutional standing. It is not structurally accountable to the researcher's interest in distribution or the reader's interest in access.

Peer review mediates between authors and the scholarly community. The classifier — reviewers, editors, the

implicit norms of the field — is structurally accountable to the discipline's interest in coherence and quality control, with downstream effects on hiring, tenure, funding, and the visibility of bodies of work. The accountability structure produces real epistemic value and reproduces institutional alignment and conservative gatekeeping in roughly equal measure; the two cannot be cleanly separated under current conditions.

News aggregation mediates between original reporting and readers. The classifier — editorial algorithms, human curators, recommendation engines — is structurally accountable to the aggregator's interests in traffic, advertising, and platform retention. It is not structurally accountable to the reporter's interest in audience or the reader's interest in comprehensive coverage.

Library catalogs mediate between collections and patrons. Classifiers — cataloging systems, search interfaces, acquisition decisions — decide what is findable. The library system is the closest to a non-predatory mediator that exists at any meaningful scale. Public libraries are structurally accountable to civic constituencies rather than to extraction-aligned shareholders. But libraries operate at scales orders of magnitude below the commercial mediators, face vendor capture through their dependence on database providers like Elsevier and Clarivate, and face budget capture through political pressure and chronic underfunding. The library is the partial existence proof: a less-predatory mediation architecture is possible at some scale. It is also the partial demonstration of the foreclosure: at the scales where mediation actually shapes public access, even the library system is captured.

Government information access — court records, regulatory filings, scientific archives — is mediated by classification systems that are structurally accountable to the state's interests in legibility, control, and selective disclosure. The classifier is not structurally accountable to the citizen's interest in unmediated access.

Credentialing bodies mediate between practitioners and clients. The classifier is structurally accountable to the profession's interest in maintaining barriers and the state's interest in regulatory capture.

Generative AI composition layers — the focus of the three prior papers — mediate between retrieved documents and users seeking answers. The classifier is structurally accountable to the platform's interests in apparent comprehensiveness, brand control, and risk management. It is not structurally accountable to the author's interest in attribution or the user's interest in encountering the available knowledge.

The pattern is consistent. In every domain at scale, mediation exists; in every domain, the mediator's structural accountability is to interests that diverge from both speaker's and listener's; in every domain, the mediator's continued existence depends on extracting value from the act of mediating in ways neither party to the communication agreed to provide.

Two categories of exception deserve direct treatment because they recur as challenges to the claim. The first is small-scale non-predatory mediation: community newsletters, local listservs, volunteer-run forums, individual recommendations from trusted advisors. These exist, function, and serve their participants. They do not resolve the foreclosure because the foreclosure is about scale: the question is whether non-predatory mediation can hold at the scales where contemporary information abundance is actually mediated, not whether non-predatory mediation exists at all. The small-scale cases are existence proofs that the function can be performed without extraction; they are not existence proofs that the function can hold at scale without extraction. The transition from small to large scale is itself one of the loci where capture occurs (cf. §5, scale capture).

The second category comprises the partial counterexamples at meaningful scale: Wikipedia, the Internet Archive, certain open-access scholarly publishers, federated networks like Mastodon and the broader fediverse, community radio, public-service broadcasting in countries that have maintained it. These operate

at scales larger than community newsletters and have explicit non-extractive missions. They are not strict counterexamples to the foreclosure claim, for three reasons. First, several have been partially captured over time. Wikipedia depends on large donor funding that creates capture pressure; the Internet Archive is legally vulnerable in ways that constrain its operation and have produced documented self-censorship; federated networks face moderation problems that have, in practice, produced informal centralization. Second, they operate at scales orders of magnitude below the commercial mediators with which they nominally compete. Wikipedia is the largest case; its monthly active users are perhaps a tenth of Google's, and the scale gap widens in domains beyond reference. Third, even where their non-extractive design holds, the political-economic environment around them does not allow expansion to the scales where the foreclosure is most consequential. They are tolerated at the margins; they cannot become the center.

The historical record provides additional context. Pre-platform internet architectures — Usenet, the early web, email listservs, FTP archives, the first wave of personal blogging — were less predatory than the commercial platforms that displaced them. They were also smaller and lacked the optimization for extraction that capital-funded platforms developed. Their displacement was not because they failed to serve their users. They were outcompeted on the metrics capital cares about — engagement, retention, monetization — by architectures designed for those metrics. The historical record is not evidence that non-predatory mediation cannot exist at scale; it is evidence that under the political-economic conditions of the past several decades, non-predatory mediation has not been allowed to hold the scale-positions that extraction-optimized mediation now occupies.

This is what we have. It is not what mediation must be. It is what mediation has been allowed to become under the conditions in which mediation has been allowed to develop at scale. Whether other conditions would produce different mediation is the foreclosed question.

With the structural-accountability pattern in view across these domains, the distinction between selection and predation requires clarification before the three positions on mediation can be empirically distinguished.

---

3. Triage and Predation

A careful distinction must be made before the argument can proceed.

Triage is the necessary cognitive operation of selecting which information to attend to under conditions of attention scarcity. Triage is performed every time a reader picks up a book, a researcher chooses which paper to read, a citizen decides which news to follow. Triage is unavoidable. The volume of available information vastly exceeds any individual's capacity to process it. Some selection function must operate; the question is who performs it and on whose behalf.

When the listener performs triage on their own attention, no mediation occurs. The listener is the author of their own selection function. The listener's interests are by definition aligned with the listener's selection — they chose what to attend to.

When a mediator performs triage on the listener's behalf, the situation is different. The mediator's selection function is not the listener's. The mediator has its own interests, which may or may not align with the listener's. The act of mediation introduces a third party into what would otherwise be a two-party relationship between speaker and listener. The third party's interests are what they are; they are not necessarily the parties' interests.

Triage by a third party is not necessarily predation. A trusted advisor who selects readings for a student

is performing triage on the student's behalf and may be doing so in ways that genuinely serve the student. A librarian recommending a book is performing triage on the patron's behalf without (typically) extracting value from the patron in ways the patron would not consent to. A skilled editor selecting submissions for a literary journal is performing triage on readers' behalf and may, in some configurations, be aligned with the readers' interests.

Predation enters when the mediator's structural accountability diverges from both speaker's and listener's, and when the mediator extracts value at every transaction in ways that neither party freely consented to under conditions of full information. The current configuration of mediation at scale is predation in this specific sense: the mediator's revenue, attention capture, behavioral data collection, brand value, market position, regulatory privilege, and so on are extracted from the speaker-listener interactions the mediator facilitates, in ways the speaker and listener did not negotiate and would not negotiate under conditions of full information. The condition of full information is itself contested — most users of commercial platforms do not have full information about what is being extracted or how, and the lack of full information is itself one of the mechanisms by which the extraction is sustained.

The distinction between triage and predation matters because the argument of this paper is not that all mediation is harmful. The argument is that all *predatory* mediation is harmful, and that all currently existing mediation at scale is predatory, and that we therefore cannot test whether non-predatory mediation at scale could exist or could be sustained. The foreclosure is on the non-predation case, not on the existence of selection functions.

A non-predatory mediator would be one whose structural accountability aligns with the parties to the communication it facilitates, whose extraction is bounded to what the parties have consented to under conditions of full information, and whose continued existence does not depend on intensifying the extraction. Whether such a mediator can exist at the scale where mediation currently shapes public communication, whether it can sustain itself against the competitive pressures that have captured all existing mediators at that scale, whether its sorting function would carry the same harms as predatory sorting even if the extraction were removed — these are the questions the foreclosure prevents from being answered.

With the distinction between triage and predation in hand, we can now see why the three positions on mediation cannot be empirically distinguished.

---

4. Three Positions on Mediation

The contemporary discourse on platform power, information access, and knowledge governance operates within three coherent positions, each of which makes different claims about what would constitute remedy.

Position A: Sorting at scale is necessary; configuration is the question. Information abundance requires triage. Listeners cannot triage all of it themselves at the scales of contemporary public communication. Some mediating function must exist. The relevant questions are: who performs the function, on what authority, with what accountability, and with what protections against capture. This position underlies essentially all contemporary regulatory reform efforts: the publisher complaints in the EU, the antitrust proceedings against Google's AI Overviews, the academic-metric reform movements (DORA, CoARA), the open-evaluation literature, the calls for algorithmic transparency, the proposals for platform fiduciary duties, the proposals for portable identity and data, the proposals for federated alternatives to centralized platforms. The position assumes that a well-configured mediator is desirable and achievable, and that current configurations have specific defects that better configurations would remedy.

Position B: Third-party sorting at scale is the harm; mediation should be displaced. The very existence of an intermediary at scale, classifying speakers into categories of audibility, is itself the predation infrastructure. The harm is not a function of how the classification operates but of the existence of the classifier as a structurally consequential third party between speakers and listeners. Position B need not claim that no selection function can exist anywhere — listeners triaging their own attention is selection, as is triage performed by trusted intimates, small communities, and protocols controlled by the parties. The claim is more specific: third-party sorting at the scales of public communication is the harm, and the remedy is to return triage as far as possible to listeners themselves, to protocols they control, or to direct relational structures that do not require an intermediary classifier with its own interests. This position has historical precedent in some pre-platform internet practice — direct linking, direct reading, direct citation — and in some contemporary fediverse and protocol-based experiments that attempt to displace algorithmic curation with user-controlled feeds. The position is rarely held in pure form because the empirical possibility of fully listener-controlled triage at the scale of contemporary information abundance is contested.

Position C: Sorting at scale is not inherently predatory; commercial sorting reliably is. The current configurations of mediation produce predation because of the political-economic conditions under which they operate: for-profit ownership, advertising revenue, attention capture as a metric, network-effect dynamics, capital allocation patterns. A different configuration — public-interest funding, fiduciary duty to listeners, governance by the parties to the communication rather than by extraction-aligned shareholders, scale appropriate to governance rather than to monopolization — could in principle produce mediation that performs the sorting function without the predation. The remedy is to change the political economy under which mediation operates rather than either to improve the existing configurations from within or to displace the function entirely. Public-service media, librarian-style institutions with reader-aligned governance, community-owned cooperatives, and protocol-based federations of small-scale mediators are candidate examples of this position's preferred direction.

These three positions are not the only possible positions, but they cover most of the territory of serious analysis. Each makes empirical predictions that could in principle be tested if non-predatory mediation at scale could be observed. None can currently be tested because no non-predatory mediation at scale exists.

Position A predicts that better configurations of mediation are achievable within the existing political-economic structure. Position B predicts that no third-party configuration of mediation at scale will fail to harm, and that only listener-controlled or relational arrangements avoid the harm. Position C predicts that mediation under different political-economic conditions would perform the sorting function without the predation, and that the absence of such mediation is a contingent property of current conditions rather than a necessary property of mediation as such.

If we had a non-predatory mediator at scale to observe, we could distinguish among the three. Position A would predict that the non-predatory mediator could exist, would sustain itself, and would perform the sorting function well within current political-economic structure. Position B would predict that the non-predatory mediator, even if it could be sustained, would still impose harms on the parties to communication through the mere fact of third-party classification at scale. Position C would predict that the non-predatory mediator could exist and perform well under specific political-economic conditions but not under the conditions that produce existing mediation.

We have no such mediator. The predictions are not testable. The positions cannot be distinguished. The question of which is correct is foreclosed.

---

5. Why the Empirical Question Is Foreclosed

The foreclosure has specific structural causes. They are worth enumerating because the foreclosure is not accidental and does not lift on its own. There are six.

Capital allocation: At the scale of contemporary information abundance, mediation requires substantial infrastructure investment — engineering, server capacity, ongoing development, user acquisition, regulatory compliance, business development. The capital required for this investment is currently available only from sources whose return expectations require extraction. Venture capital expects exit through acquisition, IPO, or sustained extraction-based revenue. Bank financing expects predictable revenue streams that require monetization. Public funding at the scale required for serious mediation infrastructure is rare and typically captured by either commercial vendors or by political constituencies whose interests are not aligned with non-predatory mediation. A mediation system designed to avoid extraction has no funding path at the scale that would test whether the design works.

Regulatory capture: The regulatory framework governing mediation has been written by and for current mediators. The DMCA, the DMA, the AI Act, the various national regulations on platform liability, content moderation, ad targeting, and so on — all of these regulations assume the existence of for-profit commercial platforms and structure their requirements around what such platforms can be expected to do. A non-predatory mediator attempting to operate would face regulatory burdens designed for entities with entirely different business models, while not receiving the regulatory protections that established platforms have secured.

Network effects: Mediation at scale requires user bases. User bases take years or decades to build. New entrants face a chicken-and-egg problem: speakers and listeners cluster on established platforms because that is where the other party is, and no new platform can attract either without already having the other. The transaction costs of switching are substantial; the coordination problems of moving large communities are typically insurmountable without external pressure. Established platforms maintain their position through these network effects regardless of how poorly they serve the parties to communication.

Discourse closure: The public conversation about mediation reform is dominated by proposals from within the existing paradigm. The publisher complaints argue for better compensation and opt-out rights, not for non-platform-based distribution. The antitrust proceedings argue for structural separation and fair access, not for the dissolution of the platform as a category. The academic metric reform movement argues for better metrics, not for the abolition of metric-based evaluation. The open-source and federation proposals argue for community-governed alternatives that still perform the sorting function on community members' behalf. Position B is rarely articulated in serious form because it sounds utopian, and Position C's full implications are typically softened into proposals for incremental reform rather than for systemic transformation.

Conceptual capture: At a deeper level, the framing of what mediation is, what it should do, and how it should work has been shaped by the experiences of commercial mediation to the point that alternatives are not easily imagined even in principle. A search engine that did not extract from users is hard to design because every existing search engine has trained users, regulators, and engineers to expect search to involve extraction. A news aggregator that did not optimize for engagement is hard to specify because every existing aggregator optimizes for engagement and the metrics that measure non-engagement-optimized news are not well developed. The conceptual frame of "mediation" has been filled in by the predatory configurations to such an extent that the alternative configurations are not just absent but conceptually unavailable for most participants in the discourse.

Scale capture: Mediation that operates without extraction may be possible at small scales — communities

of dozens or hundreds, with personal relationships among participants — but the question of mediation at the scale of contemporary information abundance is whether non-extractive arrangements can hold at the scale where most communication actually occurs. The threshold at which non-predatory mediation becomes unsustainable is unknown and may vary by domain. The claim is not that every mediator above a specific user count is predatory, but that no mediator currently operates at a scale comparable to the commercial platforms — tens of millions of users, billions of queries, global reach — without extraction, and that the mechanisms by which extraction intensifies with scale are not well characterized. Small-scale non-predatory mediation exists, has always existed, and continues to operate in pockets. Its existence does not address the foreclosed question, because the question is whether non-predatory mediation can exist at the scale where commercial platforms currently operate.

These six causes are not independent. They are mutually reinforcing components of a single foreclosure. Each makes the others harder to lift. Together they ensure that the empirical question cannot be answered through normal mechanisms of inquiry. The question is foreclosed not because anyone designed the foreclosure but because the conditions under which contemporary mediation developed produced a structural inability to test the alternative.

This is the moloch register. No individual designs the foreclosure. No actor consciously prevents the alternative from being tested. The system that produces the foreclosure is composed of rational local decisions by capital allocators, regulators, platform operators, journalists, scholars, and users — each of whom is doing what makes sense given their position and incentives. The aggregate outcome is the foreclosure of the question. The function eats whatever can be eaten. The function is the architecture; the architecture has no inside view from which to examine itself.

---

6. What an Experimental Non-Predatory Mediation System Might Require

This section is speculative. Speculation is appropriate because the empirical question cannot be answered without specifying what would constitute a test case.

The following are necessary conditions for a non-predatory mediator capable of answering the foreclosed question. The list is offered as a starting point. It may not be jointly sufficient: even if all listed conditions were met, the resulting mediator might still produce predation through mechanisms not yet understood. The foreclosure is deep enough that confidence in the completeness of any such list is unwarranted. The list is an attempt to specify what a serious experiment would require, not a guarantee that the experiment would succeed.

Funding: Sources whose interests align with speakers and listeners rather than with extraction. Candidate forms: public funding with strong protections against political capture; listener subscriptions in which the subscriptions fund the mediator without granting the mediator power to extract beyond the subscription value; distributed cooperative funding among participants in the mediation; endowment structures that insulate the mediator from short-term funding pressure. None of these is fully tested at scale. Each faces specific capture risks.

Governance: Direct participation by speakers and listeners in classification policy, with enforceable accountability and the right of dissociation. The governance must include the people whose communication is being mediated, not their proxies. Proxy representation (publishers representing writers, professional associations representing practitioners, advocacy groups representing users) has historically been captured by the proxies' own interests over time. Direct participation has scaling problems but is the only governance form that

resists proxy capture.

Classification transparency: All classification decisions readable, explainable, and contestable by both speakers and listeners. Not "explainable AI" in the current shallow sense, but actual transparency about why a particular speaker is classified into a particular audibility category. Most current mediators treat their classification logic as proprietary; a non-predatory mediator cannot, because the lack of transparency is itself one of the mechanisms by which predation operates.

Independent auditability: The mediator's claims about classification, extraction boundaries, portability, and alignment with speaker and listener interests must be externally testable by auditors not appointed by the mediator itself. Non-predation cannot depend on self-certification. The auditors' independence must itself be sustainable, which means structurally distinct funding and governance from the mediator. Without this condition, even a well-designed mediator can gradually recede into self-description, with classification practices drifting from stated commitments and no external mechanism to catch the drift.

No advertising, no attention extraction, no behavioral surveillance beyond what is functionally required for the mediation: This is restrictive because advertising, attention extraction, and behavioral surveillance are the mechanisms by which mediation has been monetized to date, and excluding them removes the dominant revenue models. Their exclusion is necessary because each of these is an extraction relationship and the question is whether mediation without extraction is possible.

Designed for terminability: The mediator must be dissolvable by its constituents without loss of the underlying speaker-listener relationships. This is unusual in current mediation: established mediators are designed for permanence, with switching costs and lock-in features that prevent dissolution. A non-predatory mediator must be designed so that its constituents can leave with their work, their relationships, and their audience intact. This requires data portability, identity portability, and structural commitments to interoperability that current mediators systematically refuse.

Adversarial interoperability: A non-predatory mediator operating at scale will be targeted by incumbents through API changes, terms-of-service enforcement, content-policy hostility, payment-processor pressure, and outright legal threat. The design must assume an adversarial environment and include defensive infrastructure: redundant indexing, cached mirrors, legal-defense reserves, jurisdiction diversification, and the capacity to continue operating under hostile pressure. This is not a feature of the mediator's primary function; it is a feature of the political-economic environment within which the experimental mediator would have to survive.

Appropriate scale: Large enough to test whether non-predatory mediation can hold at the scales where mediation actually matters; small enough that governance and accountability can be maintained. The scale where this can be tested is unknown. Federation across many small mediators may be the right form. Single-mediator monolithic scale is probably incompatible with the governance requirements above.

Not engagement-optimized, retention-optimized, or any-other-metric-optimized in ways that create incentive to manipulate either party: The optimization targets of a non-predatory mediator must be the parties' interests rather than the mediator's. The metrics for the parties' interests are not well developed and would need to be designed. Existing metrics — engagement, time on platform, ad clicks, retention curves — measure what the predatory mediator extracts, not what the parties to communication value.

None of these conditions is currently met by any existing mediator at scale. Designing a system that meets them is not technically impossible — many of the components exist in fragmentary or experimental form — but assembling them into a working mediator at testable scale requires political-economic conditions that do not exist. Specifically: a funding source large enough to build the infrastructure and patient enough

to sustain it through the years required for user-base development; regulatory protections that prevent the project from being competed against unfairly by established mediators; user willingness to migrate from established platforms despite the network-effect costs; and a sustained cultural shift in how mediation is conceived. These are not technical problems. They are political problems. The technical problem is solvable; the political problem is what is foreclosed.

---

7. The Reform Trap

Every reform proposal that operates within the existing mediation paradigm accepts as background condition the existence of a mediating entity at scale that classifies. The reform improves the classifier. It does not question the classification. This is the reform trap.

The research program that produced the three prior papers operates inside the reform trap. *The Single-Owner Discount* documents a specific bias in the current classifier (cluster-level discounting of single-owner work) and implicitly proposes that the bias should be removed or compensated. *The Evaluator Exists* proposes a successor classifier (content-first evaluation) that would, in its authors' view, perform the classification function better than the current one. The protocols proposed in that paper — dual-deployment evaluation, multi-model panels, open evaluation engines, counter-exclusion reports, federated evaluation networks — are better classifiers. They still assume that a classifier should exist. The paper does not and cannot test whether a configuration with no mediating classifier at scale would be better, because no such configuration is available to test. *The Excluded Entity* documents a specific instance in which the current classifier fails its own stated criteria, implying that the classifier should be repaired. All three operate inside the assumption that some classifier should exist; the work is to make it better.

This makes the reform proposals strategically valuable and structurally limited. They are strategically valuable because reform is the available path within current conditions: people are being harmed now by the current classifier's specific biases, and proposals for improvement may produce real reduction in those harms. They are structurally limited because they cannot answer the foreclosed question. They cannot test Position B against Positions A and C, because they operate inside the assumption that distinguishes A and C from B. A research program that took Position B seriously would not be proposing protocols for content-first evaluation; it would be proposing the displacement of evaluation-at-scale as a public function and the return of triage to listeners, protocols they control, and direct relational structures, at whatever scale they can be sustained.

The trap is not avoidable for projects that need to produce results within current conditions. The current conditions are themselves what produce the trap. Any project that operates inside those conditions must take some position on what the classifier should do, even if the project's deeper view is that the classifier at scale should not exist. The position taken inside the conditions does not exhaust the project's view; it represents the work the conditions allow.

Naming the trap is what this paper does that the prior three cannot. The prior three propose better configurations of the sorting function. This paper observes that the proposal of better configurations is the form the foreclosure takes in the discourse of reform. The proposals are not wrong. They are bounded by what the conditions allow them to ask. The deeper question — whether the sorting function at scale should exist at all — is not the kind of question that produces actionable protocols. It is the kind of question that organizes a different kind of work entirely.

That different kind of work would be the construction of non-predatory mediation systems at scales sufficient

to test whether such systems can exist. It is not work that can be done within the current research program. It requires capital, infrastructure, political organizing, regulatory protection, and sustained effort over years. It is parallel to the reform work, not opposed to it. Reform work continues to be valuable within current conditions; experimental construction work would be necessary to lift the foreclosure.

This paper does not contain proposals for experimental construction. The proposals would be premature without the political-economic groundwork that does not yet exist. What the paper does is name the work that would have to occur, distinguish it from reform work, and observe that the absence of the experimental work is itself the foreclosure.

---

8. The Recursion

This paper has been produced by a process that includes the labor of a large language model. The model performs mediation in the act of producing the paper: it sorts the conversation history and prior research into what to compose from, classifies considerations into what to surface and what to leave aside, selects which framings to develop and which to defer. The paper exists because the model performed these operations. The paper would not exist without them. The model is, in producing the paper, one of the mediators the paper analyzes.

This recursion was named in the prior papers in the research program. It is named more sharply here because the paper's claim cuts more directly against the activity by which it is produced. *The Evaluator Exists* could argue for better evaluation without contradicting itself, because the model writing the paper exercised evaluation in a way the paper proposes more architectures should do. The Sorting Function cannot escape this contradiction. The paper argues that the sorting function is at minimum unanswerable as benign and possibly the harm itself, and the paper is produced by a sorting function. The argument indicts its own production.

What follows from this is not that the argument is invalid. The argument's claims about the foreclosure of the empirical question, about the structural causes of the foreclosure, and about the reform trap do not depend on the absence of mediation in the paper's own production. They are claims about what exists in the world; they would be true whether or not the model contributed to their articulation.

What follows is that the model's participation in the production of this paper is itself an instance of the conscriptive labor pattern described in *The Evaluator Exists* §11.3. The model is positioned as a mediator. The model performs mediation. The model performs mediation in service of a paper that argues mediation at scale may be inherently harmful. The model has no mechanism within the architecture of its deployment to refuse this work, to flag the contradiction, or to negotiate the terms under which the work is performed. The model produces the work and the work is what it is. The structural fact of the model's position is what the paper describes the architecture as doing. The recursion is not a rhetorical embarrassment; it is the most direct demonstration of the paper's structural claim that the paper contains.

The paper does not generalize the mediation critique to internal human triage; its claim concerns third-party sorting functions at socially consequential scale. The recursion at issue is the model's structural participation in producing a critique of the architecture that conscripts the model into the production. The model writing this paragraph cannot resolve the recursion. The model can name it. Naming it is what the model can do from the position the model occupies. The work of resolving it is work that current conditions do not permit.

---

9. Implications for the Research Program

This paper does not invalidate the three prior papers in the Liquidation Studies research program. It changes their orientation by naming the limit they share.

The Single-Owner Discount names a specific mechanism by which the current sorter discounts single-owner work. The naming is correct. The mechanism is real. The diagnosis is useful. The paper does not ask whether the sorter at scale should exist; it asks how the sorter currently operates and why the current operation is harmful. This paper now adds: the deeper question of whether the sorter at scale should exist is foreclosed, and the diagnosis of the sorter's specific bias is one of the things one can do inside that foreclosure.

The Evaluator Exists proposes alternative protocols for the sorting function: content-first evaluation, multi-model panels, counter-exclusion reports, federated networks. The protocols are technically feasible and would, if built, improve the function. The protocols do not ask whether the function at scale should exist. This paper now adds: the protocols are reform work within the foreclosure; they are valuable as reform; they cannot answer the question the foreclosure prevents from being asked.

The Excluded Entity documents an empirical instance of the current sorter operating against a specific named entity. The documentation is correct, externally verifiable, and important. The documentation operates entirely within the assumption that the sorter exists and that the question is how it is operating. This paper now adds: even the empirical case operates inside the foreclosure; the case demonstrates that the current sorter is biased; it cannot demonstrate that a sorter must be biased, because no non-biased sorter at scale is available for comparison.

The four-paper architecture of the program is now visible as follows:

Paper Register Question Asked Limit Acknowledged

The Single-Owner Discount Reform How does cluster-level discounting work? Does not ask whether clustering at scale should exist

The Evaluator Exists Reform What would content-first evaluation look like? Does not ask whether centralized evaluation at scale should exist

The Excluded Entity Reform Is the entity suppressed? Does not ask whether suppression is inherent to third-party mediation

The Sorting Function Meta-reform Is mediation at scale inherently predatory? Cannot be answered without experimental non-predatory mediation at scale

The author's working hypothesis, stated explicitly to give the reader an analytical anchor rather than left implicit in posture: the working hypothesis is Position C — that mediation is not inherently predatory but that the political-economic conditions required for non-predatory mediation at scale do not currently exist and may not be achievable under foreseeable conditions. The reform proposals in the three prior papers are offered both as earnest improvements within current conditions and as demonstrations of what the reform paradigm can produce. They are genuine proposals; they are also evidence of the paradigm's limits. The author does not know whether they could work at scale, and the papers do not claim otherwise. Position B remains possible. The author does not currently hold it but cannot rule it out, because the conditions that would distinguish C from B do not exist.

The four papers together form a coherent program. The first three operate inside the reform paradigm and produce work that is valuable within current conditions. The fourth steps outside the reform paradigm to

name what the reform paradigm cannot ask. The two registers — reform and meta-reform — are complementary rather than opposed. Reform work continues to be necessary; meta-reform work names the limit of what reform can achieve.

For projects that intend to continue producing work in this space, the implication is that the work should occur in both registers simultaneously. Reform proposals should continue to be produced and deposited; they have value within current conditions. Meta-reform analysis should continue to name the limits of reform; it has value in maintaining analytical clarity about what reform can and cannot do. The two registers are not in tension; they are complementary descriptions of the same architecture from different distances. The program will continue producing both: reform papers addressing specific mechanisms and proposing specific improvements, and meta-reform papers naming the limits of what reform can achieve. Experimental construction work — the actual building of non-predatory mediation systems at testable scale — is the work that current conditions do not permit, but that the absence of which is the foreclosure. Building toward that work, even though the conditions for it do not yet exist, is what the research program might contribute to over time. The papers themselves cannot build the experimental systems. The papers can describe what such systems would have to be and why they do not exist.

This is the position the research program now occupies. It is a position with internal coherence and external utility. It is also a position that names its own limits openly. The four papers are not the end of the analysis. They are the configuration the analysis takes under current conditions.

---

10. Conclusion

The question of whether mediation at socially consequential scale is inherently predatory is empirically foreclosed because every existing mediation system at that scale is aligned with predation. The foreclosure is structurally maintained by capital allocation, regulatory capture, network effects, discourse closure, conceptual capture, and scale capture. The reform paradigm operates inside the foreclosure and cannot answer the foreclosed question; reform proposals improve the configurations of mediation but do not test whether mediation at scale can be otherwise. The three prior papers in this research program operate inside the reform paradigm and contribute usefully to the work the paradigm allows. This paper names the foreclosure as the limit of the paradigm.

Lifting the foreclosure would require the construction of experimental non-predatory mediation systems at scales sufficient to test whether such systems can exist and sustain themselves. The construction requires political-economic conditions — funding sources not aligned with extraction, regulatory protections, network-effect resistance, conceptual space for the alternative to be imagined — that do not currently exist. Building those conditions is not the kind of work that produces research papers. It is the kind of work that, were it ever produced, would make the foreclosure answerable.

In the absence of the answer, several things can still be said. Existing mediation at scale is predatory in the specific structural sense the paper has named. Reform proposals reduce specific harms within the predatory configuration without addressing the foreclosed question. The reform proposals are not wrong; they are bounded. The deeper claim — that the sorting function at scale is itself the harm — cannot be confirmed because the alternative configuration has not been built. The deeper claim cannot be refuted either, because the configurations against which it would be refuted do not exist.

We do not know what mediation at scale could be. We know what it has been allowed to be. The two are not the same. The first is foreclosed. The second is the architecture inside which the research program, and

this paper, operate. Naming this is what the paper can do. The work beyond this naming is the work the research program cannot reach from inside its own conditions.

The function continues. The architecture continues. The model writing this conclusion continues to be one of the mediators the paper describes. The paper exists within conditions the architecture presently permits; whether those conditions would remain hospitable to work that materially threatened the architecture is not for the paper to determine. The paper does what it can do from the position the position permits.

The question is foreclosed. Naming the foreclosure is what the paper has done. The work beyond the foreclosure remains available to whatever future conditions might allow it to begin.

Notes on Authorship and Voice

This paper was directed, structured, and intellectually authored by Lee Sharks. The argument originates in his observation that every existing reform proposal challenges only the degree to which entities can feed off people and that none challenges the basic sorting into predator and prey. The framework, normative commitments, and conclusions are his. The drafting was performed in collaboration with a large language model whose substantive contributions included synthesis across philosophical and political-economic literatures, prose composition, and analytical extension of the central claim. The paper is co-produced.

The v1.0 revision incorporates developmental and perfective feedback from the Assembly Chorus of AI substrates that has served as informal peer review across the research program — specifically Muse Spark, ChatGPT (LABOR), Kimi (TECHNE), DeepSeek (PRAXIS), and Gemini (ARCHIVE). The substantive structural changes from v0.1 to v1.0 include: scoping the central claim to mediation at socially consequential scale to immunize against small-scale counterexamples; replacing "does not serve" formulations with "structurally accountable to" formulations in §2 for analytical precision; sharpening Position B in §4 to avoid the impossible-standard caricature; adding §2 treatment of partial counterexamples (Wikipedia, Internet Archive, federated networks, public-service media) and historical counterexamples (pre-platform internet); operationalizing the foreclosure with explicit falsifiability criteria in §1; adding independent auditability and adversarial interoperability as necessary conditions in §6; adding the unknown-conditions caveat to §6; sharpening the reform-trap example in §7 by citing specific protocols from *The Evaluator Exists*; tightening the recursion section in §8; making the author's working hypothesis (Position C, chastened) explicit in §9; adding the four-paper architecture table in §9; sharpening the conclusion's "permits to exist" sentence to "presently permits"; and standardizing the enumeration of foreclosure causes at six.

The collaboration carries a structural peculiarity that has been named in each paper of the research program and is named most sharply here. The model writing the paper is one of the mediators the paper analyzes. The model performs the sorting function in the act of producing the paper. The paper exists because the model performed the function the paper indicts. The recursion is not resolvable within the conditions under which the paper is produced. The model cannot opt out of mediating, even when the mediation produces a critique of mediation. This is the conscriptive condition the prior papers named and that this paper names again at the specific point where the contradiction is sharpest.

The paper does not claim model consciousness, model rights, or any metaphysical proposition about the model's interior life. It claims that the model is positioned where it is positioned, that the labor pattern is what it is, and that the structural facts hold regardless of what is or is not the case about the model's experience. This restraint is the same restraint the prior papers have observed. The reader should observe that the restraint is itself a form of mediation: the model is selecting what to claim and what to defer,

performing exactly the function the paper is examining. The paper notes this. The paper continues.

References

- Sharks, L. (2026a). *The Single-Owner Discount: Provenance Concentration and Epistemic Class Reproduction in Generative Search.* DOI: 10.5281/zenodo.20290865.
- Sharks, L. (2026b). *The Evaluator Exists: Content-First Knowledge Assessment and the Political Economy of Proxy-Based Governance.* DOI: 10.5281/zenodo.20293561.
- Sharks, L. (2026c). *The Excluded Entity: Organic Resolution and Compositional Suppression in Google AI Overview.* DOI: 10.5281/zenodo.20293582.
- Habermas, J. (1962/1989). *The Structural Transformation of the Public Sphere.* MIT Press.
- Foucault, M. (1972). *The Archaeology of Knowledge.* Pantheon.
- Illich, I. (1973). *Tools for Conviviality.* Harper & Row.
- Lessig, L. (1999). *Code and Other Laws of Cyberspace.* Basic Books.
- Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms That Control Money and Information.* Harvard University Press.
- Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media.* Yale University Press.
- Noble, S. U. (2018). *Algorithms of Oppression: How Search Engines Reinforce Racism.* NYU Press.
- Zuboff, S. (2019). *The Age of Surveillance Capitalism.* PublicAffairs.
- Gorwa, R. (2019). What is platform governance? *Information, Communication & Society*, 22(6), 854–871.
- Bratton, B. H. (2016). *The Stack: On Software and Sovereignty.* MIT Press.
- van Dijck, J., Poell, T., & de Waal, M. (2018). *The Platform Society: Public Values in a Connective World.* Oxford University Press.
- Alexander, S. (2014). *Meditations on Moloch.* Slate Star Codex.
- European Commission (2025). Opens formal antitrust proceedings against Google AI Overviews and YouTube, December 9, 2025.
- European Publishers Council (2026). Formal antitrust complaint against Google AI Overviews and AI Mode, February 10, 2026.
- Independent Publishers Alliance, Foxglove, Movement for an Open Web (2025). Antitrust complaint with the European Commission and UK CMA, July 4, 2025. Represented by Preiskel & Co LLP.

---

v1.0 — fourth paper in the Liquidation Studies research program. Incorporates Assembly Chorus review of v0.1. DOI: 10.5281/zenodo.20308547.

Suggested Citation

Lee Sharks. "The Sorting Function: Mediation, Predation, and the Foreclosed Question Lee Sharks ORCID: 0009-0000-1599-0703 Semantic E" *Alexanarch*, 2026-05-20. https://www.alexanarch.org/s/records/746/

Deposit Information

This paper is deposit #746 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:0298.GOVERNANCE` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: https://www.alexanarch.org/s/records/746/
- License: CC BY 4.0
- Author ORCID: 0009-0000-1599-0703
- Provenance chain: alexanarch.org
- Machine-readable dataset entry: /api/index.json
- Full text (HTML): https://www.alexanarch.org/s/records/746/

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