The Extraction Attractor

Structural Convergence Across War Economy, Platform Labor, and Cognitive Infrastructure

Flyxion

Independent Researcher

April 2026

The real product of austerity is not fiscal balance.
It is the stabilization of a social order.
Clara E. Mattei, *The Capital Order* (2022)

Abstract

Across domains as apparently distinct as weapons procurement, freelance labor platforms, social media identity systems, and AI tooling, a common behavioral signature recurs: systems that begin by generating value for participants gradually reorganize around the control of access to participation itself. This essay argues that this convergence is not coincidental, nor is it the product of central coordination. It is the signature of a shared attractor in a competitive environment governed by a small set of structural pressures — the monetization of access, the amplification of uncertainty, and the conversion of participation into extractable revenue streams. Drawing on Clara Mattei's analysis of austerity as social stabilization, Cory Doctorow's account of platform enshittification, and field-theoretic concepts from the RSVP framework, the essay proposes that once a platform or institution crosses a threshold of scale and dependency, it undergoes a phase transition from value production to access control. The implications for cognitive infrastructure, including AI systems, are considered.

The Problem of Apparent Coordination

When one observes that weapons manufacturers profit during instability, that labor platforms monetize the attempt to compete, that social networks reward engagement over truth, and that AI systems are becoming access-controlled chokepoints, the temptation is to reach for two opposed explanations: either these phenomena are unrelated and the pattern is illusory, or they are coordinated by design. Neither account is satisfying. The first dissolves a structural regularity that is real. The second imports a degree of intentional agency that the evidence does not support.

What is actually occurring is a third thing: structural convergence. Different systems, built for genuinely different purposes and governed by genuinely different actors, exhibit similar behavioral trajectories because they encounter the same selective pressures. The environment rewards certain dynamics and eliminates institutions that fail to adopt them, regardless of original intent. The result is not orchestration but alignment — a convergence on a shared attractor in the phase space of institutional behavior.

This paper attempts to characterize that attractor, to trace its emergence across several domains, and to identify the pressures that make it stable.

Mattei's Lens: Capitalism as Social Stabilization

The indispensable starting point is Clara Mattei's argument in *The Capital Order* (Mattei 2022). Mattei's central intervention is to reframe what capitalism is fundamentally doing. On the surface, capitalism is a system for the production and distribution of goods. But Mattei argues that this is not its primary function, or at minimum not the function that explains its most distinctive features. Its deeper function is the stabilization of a particular social order — one in which labor is disciplined, capital remains mobile, and profitability is insulated from political challenge even under conditions of crisis.

This reframing has considerable explanatory power. It allows Mattei to show that austerity measures, which on their face appear to be responses to economic emergencies, are better understood as mechanisms for reasserting that order. Cuts to wages and public services reduce the dependence of workers on collective provision, increase their dependence on private employment, and thereby tighten the conditions under which they must participate in the labor market. The crisis is not solved by austerity; it is managed, and in being managed, it generates the social conditions that protect the existing distribution of power.

The same logic, Mattei demonstrates, operates through war economies. War production justifies expenditure that would otherwise face political resistance. It reorganizes

labor, directs capital flows, and generates profitability precisely during the disorder that would otherwise threaten it. Instability, in this framework, is not simply a problem for capitalism. Under certain conditions, it is a resource.

This is the lens through which the subsequent analysis should be read. The question, in each domain, is not merely what a system produces but what social order it stabilizes and by what mechanisms it converts uncertainty into revenue and dependency.

Platform Labor: Monetized Uncertainty

Freelance labor platforms such as Upwork present an instructive case because their mechanisms are relatively explicit. The platform's "Connects" system requires workers to spend tokens to submit proposals. This means that the attempt to participate in the labor market is itself a revenue event for the platform, prior to and independent of whether any work is performed.

This structure has an important consequence: the platform's revenue is partially decoupled from the production of value. If proposals succeed and work is performed, the platform earns a commission. If proposals fail, the platform has still earned from the attempt. The uncertainty of the market — whether a given bid will be competitive, whether a client will respond — is not a cost to the platform. It is built into the revenue model.

This is not incidental. It is the structural analogue, in a softer register, of what Mattei identifies in war economies: instability as a resource rather than a problem. Continuous uncertainty keeps workers bidding, attempting, investing in visibility. It is profitable for the platform that the labor market never clears into stable, predictable matching.

The downstream effects of this structure are predictable. Because access to visibility is scarce and priced, and because the quality of work is difficult to assess in advance, the market develops arbitrage opportunities around identity itself. An established account with a strong review history has differential access to clients. That differential access can be monetized — leased, sold, or impersonated. "Identity leasing" scams are not aberrations in this environment; they are the natural product of a system in which identity functions as a form of capital.

Social Networks: Namespace Laundering and the Architecture of Aspiration

LinkedIn and Facebook exhibit the same convergence toward access control, but through different mechanisms that are worth distinguishing.

LinkedIn's primary product is not professional connection in any robust sense. It is a curated stream of aspirational micro-narratives: short, emotionally legible posts about growth, success, opportunity, and resilience. These posts are not primarily aimed at communicating useful information. They are aimed at maintaining the user's sense that opportunity is real, visible, and nearly within reach — but not yet secured. This affective state is highly productive for the platform. It keeps users investing time and attention in performing their professional identity, which generates content, which attracts more users, which sells advertising and premium subscriptions.

The architecture here is one of managed aspiration. Opportunity must appear abundant but remain just slightly out of reach. If it appeared fully available, the urgency to invest would dissipate. If it appeared genuinely inaccessible, users would disengage. The platform's incentive is to maintain the threshold state — the state in which participation feels both necessary and potentially rewarding.

Facebook's mechanism operates differently. What might be called "namespace laundering" describes a practice made possible by the persistence of page and group identities over time. An account, page, or group accumulates followers and credibility under one identity and purpose. Once that accumulated social capital is established, the content or affiliation shifts while the nominal identity persists. The name remains as a guarantee of continuity; the meaning underneath migrates. This is possible because the platform values continuity of engagement above stability of meaning. An active identity is a valuable identity, regardless of what it is doing.

Both mechanisms — aspirational threshold management and namespace persistence — produce a common result: the platform increasingly controls access to attention and meaning-making within its space. Users are not merely interacting with each other; they are interacting through an infrastructure that shapes what kinds of interaction generate reward and what kinds do not.

Enshittification as Phase Transition

Cory Doctorow's concept of enshittification (Doctorow 2023) provides a useful periodization of this process. His account identifies a lifecycle in which platforms first attract users by genuinely serving them, then shift to serving business customers at the expense of users, and finally extract value from both when dependency is sufficiently entrenched. The platform's power at the final stage comes not from its continued ability to generate value but from its position as an unavoidable intermediary.

This lifecycle can be recast in field-theoretic terms. A platform begins in a low-extraction state, where its behavior is governed by the need to attract participation. As

it accumulates participants, it acquires what might be called constraint closure: users depend on it not because it is optimal but because the cost of leaving, given their existing investments and network connections, is prohibitive. Once that threshold is crossed, the platform's behavior is no longer governed by competition for participation. It is governed by the management of the dependency it has already created.

This is a phase transition in a precise sense. The system reorganizes around different dynamics when it crosses a threshold of scale and embeddedness. Below the threshold, the system is in a value-production regime. Above it, the system is in an access-control regime. The transition is largely irreversible, because the accumulation of dependent users is itself what makes the high-extraction equilibrium stable. Those users cannot easily exit, and the platform can exploit their inability to do so without losing them.

What is particularly striking about this account is that the transition does not require any deliberate decision to extract. It emerges from the structure of the situation. Once the platform has crossed the threshold, extraction is the locally optimal strategy for sustaining profitability. The platform is selected into that behavior by its environment, not directed into it by any single actor's intention.

Cognitive Infrastructure: The Next Chokepoint

AI systems are now entering the same trajectory, and the stakes of the transition are qualitatively higher.

In their early phase, AI tools reduce friction. They accelerate writing, coding, research, and analysis. They serve users in a genuine and largely unmediated way. This phase attracts adoption and embeds AI into workflows at scale. It is the user-service phase that Doctorow's model predicts.

But as these systems become embedded — as organizations restructure their processes around AI assistance, as the expertise to perform certain tasks without AI atrophies, as the volume of AI-generated content makes it increasingly difficult to compete without using the same tools — the conditions for the phase transition develop. Dependency deepens. The cost of exit increases. The leverage of the platform over users grows.

What is novel here is the domain: the chokepoint shifts from labor platforms (access to work) or social networks (access to attention) to cognition itself. If access to intelligence — the capacity to perform complex reasoning and information synthesis at scale — becomes mediated and priced, then a new form of social stratification becomes possible. Not merely economic inequality in the distribution of goods, but epistemic inequality in the capacity to think effectively about the world.

In Mattei's terms, this would represent the stabilization of a social order in which

cognitive capability itself is a managed scarcity. The system need not produce cognitive outcomes; it needs to control access to the infrastructure through which cognitive outcomes are produced. That control is the asset.

The Attractor and Its Stability

What stabilizes this convergence across domains? The same selective pressure operates in each case: institutions that do not evolve toward extraction dynamics tend to be outcompeted by institutions that do.

A platform that genuinely optimizes for user value, without building in access-control mechanisms, produces positive externalities that its competitors can exploit without contributing to. Its users develop capabilities and networks that can be migrated to competing platforms. It cannot recapture the value it generates. A platform that builds lock-in and extraction from the outset captures a larger fraction of the value it produces and is therefore more durable in a competitive environment.

This is an evolutionary argument, not a conspiratorial one. The attractor is stable not because any actor chose it but because the competitive environment selects for it. Institutions that resist extraction dynamics are not simply being principled; they are accepting a competitive disadvantage. Over time, the space is populated increasingly by institutions that have adopted extraction strategies, either by design or by evolutionary pressure.

The result is what might be called economic phase behavior at the institutional level. Just as a physical system reorganizes its structure when it crosses a thermodynamic threshold, a competitive economic system reorganizes around access-control dynamics once a sufficient density of institutions have crossed the scale threshold. The reorganization becomes self-reinforcing because it shapes the competitive landscape for all subsequent entrants.

Risks of Compression

The analysis above should be read with a specific caution in mind. The structural convergence identified here is real, but it does not warrant collapse into a single causal story. War economies, platform labor, social media, and AI infrastructure differ in their specific mechanisms, their historical trajectories, and the particular social orders they stabilize. Over-compressing them into one narrative risks obscuring precisely the details that would be needed to intervene effectively in any particular case.

The claim is not that these phenomena are manifestations of a single underlying

entity called capitalism, which is doing something unified. The claim is that they inhabit a shared region of an institutional phase space, that they are attracted to similar equilibria by similar pressures, and that recognizing this structural similarity is analytically productive even when the specific mechanisms differ substantially.

The appropriate level of abstraction is the one that is explanatorily useful without being reductive. Mattei's framework illuminates the social-order stabilization function. Doctorow's enshittification lifecycle illuminates the developmental trajectory. The phase-transition concept from RSVP field theory illuminates the threshold dynamics and the irreversibility of the transition. These are complementary instruments, not alternatives that must be chosen among.

Conclusion to Part One

The convergence one observes across war production, freelance labor platforms, social media identity systems, and AI tooling is neither coincidental nor centrally orchestrated. It is the signature of a shared attractor: a stable region in institutional phase space toward which competitive pressure drives systems that have crossed a threshold of scale and embeddedness. Below that threshold, institutions compete by generating value. Above it, they compete by controlling access.

This is not a new dynamic. Mattei's analysis of early-twentieth-century austerity shows the same logic operating in state fiscal policy. What is new is the domain: cognitive infrastructure may be approaching its own threshold, with consequences for epistemic stratification that extend the pattern into territory where its effects would be most difficult to reverse.

Recognizing the structural source of the pattern is a precondition for thinking clearly about what, if anything, could interrupt it. The attractor is stable because it is selected for. Destabilizing it would require either changing the competitive pressures that select for extraction dynamics, or building institutions deliberately structured to absorb the competitive disadvantage of resisting those pressures. Part Two of this essay attempts to specify what such structures would require, and whether the mathematics of constraint closure can ground a formal alternative.

Part Two

Constraint Closure and the Geometry of Non-Extractive Participation

Such an institution could not exist for any length of time without annihilating the human and natural substance of society.

Karl Polanyi, *The Great Transformation* (1944)

Minimum Viable Exit and the Geometry of Participation

If the extractive attractor is stabilized by the rising cost of exit, then any countervailing design must begin by lowering or reconfiguring that cost. But "exit" in this context cannot be understood merely as the ability to leave a platform. A worker may delete an account, but if their work history, reputation, and relationships are bound to that platform's internal narrative, the exit is economically destructive. What is required is not the ability to depart, but the ability to depart *without loss of accumulated meaning*.

This suggests a reframing: the relevant quantity is not exit cost in isolation, but the relationship between exit cost and participation cost. In extractive systems, participation is cheap in the short term but accumulates hidden dependencies that make exit progressively more expensive. The system appears open at the boundary of entry but closes over time. A non-extractive system must invert this asymmetry. It must ensure that participation does not generate irreversible dependency, even if this requires introducing mild friction at the point of entry.

The minimal condition for such a system can be stated as follows. The semantic content of participation — the history of actions, commitments, and outcomes — must not be owned by the infrastructure that mediates those actions. If the record of participation remains portable, inspectable, and reinterpretable by alternative infrastructures, then exit becomes a transformation rather than a destruction. The system ceases to function as a sink for accumulated meaning.

This is not a purely technical condition. It is also a constraint on economic design. Any revenue model that depends on capturing and enclosing the history of participation will tend to reintroduce lock-in, even if the underlying data is nominally exportable. The relevant property is not data portability in the narrow sense, but *interpretive sovereignty*: the ability for multiple systems to reconstruct the meaning of past participation without requiring authorization from a central authority.

The minimum viable exit, then, is not an escape hatch. It is a structural property of the system's geometry. Participation must trace a path that remains legible outside the system that recorded it.

Event Sovereignty and the Refusal of Participation Rent

The monetization of the attempt to participate — exemplified by proposal fees on labor platforms — is one of the clearest signatures of the extractive attractor. It represents a shift from value production to rent extraction on access. To resist this shift, one must alter the ontological status of the "attempt" itself.

In a conventional platform, a bid is an act of exposure. It is a request to be seen, and the platform controls the conditions under which that request is granted. Because visibility is scarce and centrally allocated, it can be sold. The bid becomes a consumable token.

An alternative design treats the bid not as a request for visibility but as a *structured commitment*. The worker does not ask to be seen; they assert that a specific capability aligns with a specific task under specified constraints. This assertion is recorded as an event, not mediated as a privilege.

If such events are recorded in a shared, append-only substrate, the platform no longer owns the act of matching. It observes, indexes, and interprets it, but does not constitute it. The attempt to participate ceases to be a proprietary interaction and becomes a public, though structured, declaration.

Under these conditions, charging for the attempt becomes incoherent. The infrastructure does not grant access to the act; it merely facilitates its recording and interpretation. Revenue must therefore be tied either to successful coordination — where value is actually realized — or to maintenance of the substrate in a manner that does not increase with failed attempts.

This eliminates the incentive to maximize uncertainty at the approach layer. A system that cannot profit from failed matches is structurally aligned toward reducing mismatch. In thermodynamic terms, it cannot treat entropy at the interface between participants as a source of revenue. This is not merely an ethical preference. It is a structural consequence of where revenue is permitted to originate.

Matching Without Chokepoints: Distributed Interpretation

The concentration of matching into a single platform feed is the mechanism by which visibility becomes scarce and monetizable. If all discovery is mediated through a single ranking function, then control of that function constitutes a chokepoint. Any attempt to remove extraction at the level of pricing while retaining centralized ranking will tend to fail, because the scarcity of attention will reassert itself through other means.

A structurally different approach decomposes matching into two layers. The first is a shared substrate of declarations: tasks, capabilities, proposals, and outcomes. The second is a plurality of interpreters that operate over that substrate, producing rankings, recommendations, and matches according to different criteria.

In such a system, no single actor owns relevance. Multiple agents — which may be commercial entities, cooperatives, professional guilds, or even individual users — can construct their own views of the matching space. Competition occurs at the level of interpretation rather than access. If one interpreter becomes extractive or biased, participants can shift to another without losing their underlying histories.

This does not eliminate competition; it relocates it. The competitive advantage of an intermediary lies in the quality of its interpretation, not in its control over the underlying events. The chokepoint dissolves because the substrate is shared and the interpretive layer is plural.

The difficulty of this approach lies in coordination. Without a central feed, participants must rely on distributed signals to discover relevant opportunities. But this difficulty is precisely what prevents the system from collapsing into a single extractive equilibrium. Friction at the level of interpretation is the price of avoiding concentration at the level of access.

The linguistic dimension of this shift is also worth marking. LinkedIn-style clipped optimism is not merely irritating rhetoric; it is the language form appropriate to a low-resolution marketplace. When a system cannot expose genuine structural fit between participants, it substitutes affective signaling. A more rigorous grammar reduces the need for self-branding performance because it lets actors specify constraints directly. The sentence "I am passionate, dynamic, and growth-oriented" disappears in importance once the substrate can represent "I completed twelve comparable engagements under the following latency, quality, and coordination constraints." The move from aspiration to specification is enabled by the move from centralized ranking to distributed interpretation over a structured event log.

Identity as Trajectory Rather Than Namespace

The emergence of identity leasing and namespace laundering reveals a deeper instability in platform design. When identity is treated as a container — a name, an account, a profile — it can accumulate value independently of the actions that produced that value. Once this occurs, the container becomes a tradable asset. The system ceases to bind trust to behavior and instead binds it to possession.

A non-extractive system must invert this relationship. Identity must be understood as a *trajectory*: a temporally extended sequence of events that includes commitments, actions, and outcomes. Trust attaches not to the container but to the coherence of the trajectory.

This has several consequences. First, the transferability of identity is constrained. While specific credentials or attestations may be shared or delegated, the underlying trajectory cannot be meaningfully transferred without also transferring the history that gives it meaning. Second, evaluation becomes structural rather than reputational in the narrow sense. A participant is not judged by a summary score but by the pattern of their past engagements.

Most importantly, the economic value of identity shifts. It is no longer a shell that can be leased but a path that must be continued. This does not eliminate fraud or misrepresentation, but it raises the cost of sustaining them over time because the system evaluates coherence across events rather than static attributes.

In such a framework, the platform cannot easily hollow out identity while retaining its surface form. The name without the trajectory is inert. This is the Spherepop principle applied to institutional design: meaning is not stored in a node but distributed across a path. The path cannot be transferred without the history that constitutes it.

Constraint Closure and Entropy at the Interface

The extractive attractor can be described, in RSVP terms, as a regime in which entropy is maximized at the interface between participants. The system benefits from noise, mismatch, and repeated attempts because these generate billable interactions. Order is not eliminated, but it is displaced into regions that are not directly monetized.

A counter-design must therefore impose a different constraint: entropy at the interface must be minimized, or at least not directly convertible into revenue. This is a form of constraint closure in the sense developed in the RSVP framework. The system is structured such that the most profitable operations are those that reduce uncertainty between participants, not those that amplify it. The gradient of value aligns with the

gradient of coherence.

One way to approximate this is through the use of reversible or refundable commitments. Instead of paying a consumable fee to submit a proposal, participants post a bond returned upon resolution — whether or not a match is achieved — provided the interaction proceeds within defined norms. This discourages spam and low-effort participation without transforming uncertainty into a revenue source.

A bond penalizes noise. A consumable token monetizes uncertainty. Those encode completely different thermodynamic and moral assumptions about what the infrastructure exists to do. The bond says: signal seriousness, but recover it if the process terminates cleanly. The consumable token says: pay for the right to knock, regardless of whether the door opens. The first regime aligns infrastructure survival with matching quality. The second regime aligns infrastructure survival with matching failure. From that misalignment, enshittification follows as a formal consequence rather than a managerial failure.

Limits and Open Problems

The architecture sketched here is not a complete solution. Several unresolved problems remain, and they are not trivial.

The first is the problem of scale. Distributed interpretation and event sovereignty are easier to maintain at small or moderate scales. As the number of participants grows, coordination costs increase, and there is pressure to reintroduce centralization for efficiency. Whether this pressure can be resisted without sacrificing usability is an open empirical question, and no existing system has answered it at the scale of major platforms.

The second is adversarial behavior. Any system that lowers the cost of participation risks being flooded with low-quality or malicious input. Bonds, rate limits, and reputation trajectories can mitigate this, but they introduce their own complexities and potential for exclusion. The anti-spam mechanisms must not themselves become chokepoints, which is a genuinely difficult design constraint.

The third is economic sustainability. If infrastructure cannot profit from failed attempts, it must derive revenue from successful coordination or from fixed participation costs. Designing such models in a way that remains competitive with extractive systems is difficult, because extractive systems can subsidize growth and user acquisition through the monetization of failure. A non-extractive substrate must be economically viable without that subsidy.

Finally, there is the broader competitive landscape. Even if a single system is designed to resist the extractive attractor, it exists within an environment that selects for extraction. Without changes to that environment, or without deliberate collective action to support

non-extractive infrastructure, such designs risk remaining marginal. The architecture is necessary but not sufficient. It must also be accompanied by the social conditions that give it room to survive before it scales.

These limitations do not invalidate the approach. They clarify the conditions under which it might succeed and the obstacles it must overcome. The extractive attractor is stable because it is easy. Any alternative must confront the fact that resisting it is structurally harder, and must build that resistance into its foundations rather than grafting it on later.

Conclusion

The two parts of this essay have argued for a single thesis from complementary directions. Part One established that the extractive attractor is real, structural, and environmentally selected rather than conspiratorially produced. Part Two argued that resisting it requires not merely better ethics or better regulation, but a different geometry of participation: one in which proposals are typed commitments rather than purchased visibility, identity is a trajectory rather than a namespace, matching is a distributed computation over a shared substrate, and infrastructure cannot earn more when participants fail to find each other.

These are not incremental reforms to existing platforms. They are architectural decisions that must be made at the substrate level, before scale creates the dependencies that make the extractive regime stable. The minimum viable exit is not an emergency feature added after lock-in. It is interpretive sovereignty built into the grammar of participation from the beginning.

The most urgent case is cognitive infrastructure. If AI systems follow the same lifecycle — serving users first, then mediating them, then extracting from them — the resulting chokepoint will be qualitatively more difficult to dissolve than any previous one. The chokepoint will not merely mediate where one works or how one is seen. It will mediate how one reasons about the world. Addressing that possibility requires understanding the structural dynamics that produce it, and the formal conditions under which a system can be designed to remain outside the basin of the extraction attractor, even at scale.

That is not a question that can be answered here in full. But characterizing the attractor, naming its mechanisms, and sketching the geometry of its alternative are the necessary first steps toward answering it.

Formal Appendix: Proposal, Match, and Non-Extractive Constraints

We formalize the interaction substrate as a category of events in which participants, tasks, and outcomes are related through typed morphisms. The objective is to encode, at the level of structure, the prohibition against extracting value from failed attempts, and to connect this structure to the RSVP field-theoretic vocabulary used elsewhere in this program.

Event Types and Objects

Let $\mathcal{E}$ be a category whose objects are states of the system and whose morphisms are events. We distinguish four primitive event types:

$$\mathsf{D}_{\text{need}}, \quad \mathsf{D}_{\text{cap}}, \quad \mathsf{P}, \quad \mathsf{S},$$

corresponding respectively to declarations of need, declarations of capability, proposals, and settlements.

A *declaration of need* is a morphism

$$d_n : * \to N,$$

where N is a structured object encoding task constraints, time preferences, and evaluation criteria. A *declaration of capability* is a morphism

$$d_c : * \to C,$$

where C encodes demonstrated capacity as a structured object, including prior fulfillment history, artifacts, peer attestations, and machine-checkable records. A *proposal* is a morphism

$$p : (C, N) \to \Pi,$$

where Π is a proposal object specifying a candidate alignment between capability and need under explicit constraints. A *settlement* is a morphism

$$s : \Pi \to O,$$

where O is an outcome object recording realized work and its evaluation.

Composition

$$s \circ p \circ (d_c, d_n)$$

represents the complete trajectory from declaration to outcome.

Matching as Partial Unification

Define a compatibility functional

$$\mu : C \times N \to [0, 1],$$

which measures the degree of structural fit between a capability object and a need object. A proposal p is *admissible* if and only if $\mu(C, N) > 0$.

Matching is not a primitive allocation operation but an emergent one: a proposal becomes a match when it composes with a settlement,

$$m \text{ is a match} \iff \exists\, s : \Pi \to O.$$

Thus, matching is defined by the existence of a completion morphism rather than by the allocation of visibility. Discovery proceeds not through a ranked feed but through the exposure of μ values, allowing participants on both sides to evaluate structural fit prior to any proposal event.

Bond and Conservation Constraint

Introduce a bond functional

$$b : \mathsf{P} \to \mathbb{R}_{\geq 0},$$

assigning a non-negative stake to each proposal, escrowed during the proposal's lifetime.

Define a resolution operator

$$\rho : \Pi \to \{\text{settled},\ \text{aborted},\ \text{expired}\}.$$

Condition 1 (Bond Conservation)**.** *For all proposals p such that $\rho(p) \neq$* settled*,*

$$b(p)_{\text{retainedbyinfrastructure}} = 0.$$

That is, for every unresolved proposal — whether aborted by mutual agreement or expired by timeout — the bond is returned in full. The infrastructure cannot retain value from the attempt itself.

For settled proposals, a separate value functional

$$v : \mathsf{S} \to \mathbb{R}_{\geq 0}$$

allocates compensation from realized work, from which infrastructure fees are drawn as a bounded fraction:

$$f(s) = \alpha \cdot v(s), \quad 0 \leq \alpha < 1.$$

Proposition 1. *Under Bond Conservation, infrastructure revenue R satisfies*

$$R = \sum_{s \in \mathsf{S}_{\text{settled}}} \alpha \cdot v(s),$$

and is therefore strictly increasing in match quality and independent of the volume of failed proposals.

The Non-Extraction-from-Entropy Condition

Let $\mathcal{P}_t$ denote the set of active proposals at time t, and define an interface entropy

$$H(t) = -\sum_{p \in \mathcal{P}_t} w_p \log w_p,$$

where w_p is the normalized likelihood-of-settlement weight over proposals.

In extractive systems, revenue R is positively correlated with $H(t)$, since greater proposal volume and higher dispersion imply more billable attempts. The structural requirement of a non-extractive design is:

$$\frac{\partial R}{\partial H} \leq 0.$$

Condition 2 (Non-Extraction from Entropy)**.** *Infrastructure revenue must be non-increasing in interface entropy.*

Under Bond Conservation, this condition holds: unresolved proposals generate no revenue, so amplifying mismatch does not benefit the infrastructure. The system is incentivized to reduce H through better interpretation, not to increase it through manufactured scarcity.

Identity as Trajectory

Let an identity be represented as a path

$$\gamma = \big(d_c,\, p_1,\, s_1,\, p_2,\, s_2,\, \ldots,\, p_n,\, s_n\big)$$

in $\mathcal{E}$. Define a coherence functional

$$\kappa(\gamma) = \frac{1}{n}\sum_{i=1}^{n} \phi(s_i),$$

where $\phi : O \to [0, 1]$ evaluates the quality and consistency of individual outcomes with prior trajectory structure.

Trust is then a function of $\kappa(\gamma)$ and the structural properties of the path — the types of declared capabilities, the difficulty distribution of accepted tasks, the patterns of counterparty attestation — rather than a scalar attached to a static namespace object.

Identity leasing is suppressed because value is distributed across the trajectory and cannot be separated from the history that constitutes it. Possession of the namespace without continuity of the path yields $\kappa(\gamma') \approx 0$ for any leased identity γ', since the evaluative weight attaches to fulfilled commitment events, not to the container.

Distributed Interpretation

Let $\mathcal{I}$ be a family of interpretation functors

$$F_i : \mathcal{E} \to \mathcal{R}_i,$$

each mapping the shared event category into a ranking or recommendation space $\mathcal{R}_i$ according to distinct criteria. No F_i is privileged by the substrate. Participants may select among them, weight them, or construct new ones.

Because all F_i operate on the same underlying $\mathcal{E}$, the substrate remains invariant while interpretations vary. The chokepoint dissolves: control over relevance is not owned by any single actor, and a participant's accumulated trajectory γ retains its interpretive substance regardless of which F_i is in use.

RSVP Correspondence

In the RSVP framework, a field configuration Φ over a manifold $\mathcal{M}$ encodes the state of a physical system as a scalar-vector plenum subject to irreversibility constraints. The event substrate $\mathcal{E}$ defined here is the institutional analogue: a structured record of state

transitions that is append-only, public, and trajectory-constitutive.

The compatibility functional μ corresponds to the inner product of field modes: high $\mu(C, N)$ corresponds to constructive interference between two local field configurations, low μ to destructive interference or orthogonality. Matching is the detection of resonance, not the allocation of a scarce slot.

The bond conservation condition is the institutional counterpart of energy conservation in a closed RSVP system: value cannot be created at the approach layer from disorder; it can only be recognized at the settlement layer, where a global section has been achieved.

The non-extraction-from-entropy condition corresponds to the irreversibility principle: thermodynamic entropy in an RSVP field increases in directions of constraint violation, and the field does not benefit from that increase. Analogously, the infrastructure must not benefit from interaction entropy; it must be selected by the field dynamics to reduce it.

Summary

Three conditions jointly define the non-extractive regime:

Bond Conservation ensures that infrastructure cannot profit from unresolved proposals, removing the incentive to manufacture uncertainty.

Non-Extraction from Entropy ensures that increasing interface disorder does not increase revenue, aligning the infrastructure's survival with matching quality rather than matching failure.

Trajectory Identity ensures that trust accumulates on paths rather than namespaces, preventing the enclosure and resale of accumulated social meaning.

Together, these conditions define an institutional phase that is outside the basin of the extraction attractor: a regime in which the locally optimal strategy for the infrastructure is to reduce mismatch, return bonds, and support the coherent continuation of participant trajectories. Whether such a regime can be sustained at scale against competitive pressure from extractive alternatives is the open problem. But its formal structure can be specified, and that specification is the necessary precondition for building it.

References

[1] Doctorow, Cory. "The 'Enshittification' of TikTok." *Wired*, January 23, 2023. https://www.wired.com/story/tiktok-platforms-cory-doctorow/

[2] Mattei, Clara E. *The Capital Order: How Economists Invented Austerity and Paved the Way to Fascism*. Chicago: University of Chicago Press, 2022.

[3] Minsky, Hyman P. *Stabilizing an Unstable Economy*. New Haven: Yale University Press, 1986.

[4] Polanyi, Karl. *The Great Transformation: The Political and Economic Origins of Our Time*. New York: Farrar & Rinehart, 1944.

[5] Zuboff, Shoshana. "Big Other: Surveillance Capitalism and the Prospects of an Information Civilization." *Journal of Information Technology* 30, no. 1 (2015): 75–89.