

Optimization Without Restraint

Gradient Amplification and Extractive Convergence in Digital Platform
Systems

Flyxion

February 13, 2026

Abstract

The contemporary internet exhibits a structural drift from participatory infrastructure toward extractive platform regimes characterized by surveillance, lock-in, and asymmetrical optimization. Economic accounts such as Tim Wu's theory of extraction and Cory Doctorow's model of enshittification describe the lifecycle through which digital platforms transition from value creation to rent capture. While analytically powerful, these frameworks remain primarily within political economy.

This essay develops a structural interpretation grounded in constraint-first ontology and thermodynamic governance. Platforms are modeled as gradient-harvesting systems that convert informational asymmetries into economic work. Surveillance reduces uncertainty, optimization exploits resulting gradients, and weakened institutional constraints permit convergence toward extractive equilibria. Informational trust collapse, wage discrimination, dynamic pricing, and attention capture are treated as manifestations of unconstrained gradient descent in socio-technical systems.

Drawing on scalar–vector field analysis and entropy-bounded semantic constraint, the essay advances a constructive program centered on interoperability, reversibility, auditability, and bounded surveillance. The decay of the internet is thus understood not as technological inevitability but as constraint erosion. Restoring invariant structural limits can stabilize informational ecosystems and realign digital infrastructures with autonomy, epistemic integrity, and democratic resilience.

1 Introduction: From Extensible Commons to Extractive Regimes

The early public internet was not merely a technological substrate but an extensible environment of agency. Its architecture permitted modification, interoperability, and recoverability. Users experienced tools as predictable instruments rather than as adversarial mediators. Protocols were open, clients were interchangeable, and failure was typically local rather than systemic. The web browser could be modified; email clients could be replaced; servers could be self-hosted. Ranking systems and search outputs functioned, at least aspirationally, as epistemic aids rather than monetized influence surfaces.

This environment was neither utopian nor free of hierarchy, yet it possessed a structural property that has since eroded: reversibility. When a service degraded, alternative clients or competing platforms could arise. When an interface behaved in ways adverse to user interests, it could be modified, filtered, or bypassed. The architecture preserved degrees of freedom. Agency was not absolute, but it was structurally encoded.

Contemporary platforms differ in kind rather than merely in scale. The dominant digital infrastructures of the present moment operate as vertically integrated optimization systems. Interfaces simulate neutrality while embedding layered ranking logic tuned to advertiser yield, behavioral prediction, and engagement maximization. Search results are blended with sponsored placement; recommendation systems interleave organic and paid signals; marketplace listings incorporate fee-weighted visibility. The surface remains smooth while the objective function shifts.

This transformation is institutional as much as technical. Digital infrastructure has migrated from protocol-based commons toward proprietary, enclosure-based ecosystems. Anti-circumvention law constrains modification. Intellectual property regimes prohibit reverse engineering. Acquisitions suppress emergent competitors before they mature into structural alternatives. Disclosure-based privacy frameworks shift cognitive burden to individuals while leaving surveillance architecture intact. The result is not merely corporate concentration but objective misalignment: systems that once mediated communication now optimize extraction.

The phenomenological dimension of this shift is widely recognized. Users report diminished trust in search results and recommendation feeds. Sellers report declining margins despite rising transaction volumes. Workers encounter intensified surveillance and algorithmic management. What appears as frictionless convenience masks layered optimization targeting attention, vulnerability, and dependency.

Recent public discourse has articulated this drift through economic and lifecycle models. These frameworks identify monopoly power, bottleneck control, and lock-in dynamics as central drivers. Yet even where these diagnoses are accurate, they remain incomplete without a deeper account of why such dynamics converge so reliably under contemporary institutional conditions.

To understand the decay of the contemporary internet requires moving beneath market structure toward structural ontology. Optimization is not inherently pathological; it becomes corrosive when unbounded by invariant constraints. Surveillance does not inevitably produce exploitation; it does so when informational gradients are permitted to accumulate without reciprocal balancing. Extraction is not an anomaly but the predictable attractor of unconstrained gradient descent within socio-technical systems.

The sections that follow reconstruct the prevailing economic explanations before situating them within a broader thermodynamic and constraint-first framework. The aim is not to displace existing critiques but to embed them within a structural theory capable of generating constructive remedies.

2 Economic Diagnoses: Extraction and Lifecycle Degradation

Two influential frameworks have emerged to describe the present condition of digital platforms. The first, articulated by Tim Wu, centers on the concept of extraction. The second, advanced by Cory Doctorow, describes a staged lifecycle of platform degradation termed enshittification. Though developed independently, these models converge upon a shared structural insight: dominant platforms have shifted from value creation toward rent capture under conditions of weakened competitive discipline.

Wu’s theory of extraction draws from classical monopoly economics. Extraction occurs when a firm possessing durable market power is able to take wealth in excess of the competitive value of its service by exploiting bottlenecks, dependency, and asymmetrical information. The focus is not merely on high prices but on the systematic conversion of structural leverage into economic surplus. Under competitive conditions, firms are disciplined by exit options; under monopoly or oligopoly, this discipline erodes. The resulting environment encourages a shift from product improvement toward identification of “pain points” where consumers or suppliers lack viable alternatives.

Digital platforms provide a particularly potent substrate for extraction because they function as intermediated marketplaces. They bring together buyers and sellers,

speakers and audiences, workers and employers. When such platforms achieve sufficient scale, they acquire the capacity to modulate visibility, ranking, and access. Sponsored placement, dynamic pricing, and fee layering become mechanisms through which the platform monetizes its gatekeeping role. The system remains functional in appearance, yet the informational order within it becomes a revenue surface.

Doctorow's model complements this analysis by describing the temporal dynamics of platform degradation. According to his schema, platforms pass through three phases. In the first phase, they subsidize users to build scale and network effects. In the second, once user lock-in is established, they degrade user experience in order to attract and privilege business customers, such as advertisers or sellers. In the third phase, once both users and business customers are dependent, the platform extracts from both sides, reducing overall value to the minimum required to prevent mass exit while redirecting surplus toward shareholders.

The strength of the enshittification model lies in its attention to sequencing. It identifies not merely the presence of extraction but the structural transition from subsidization to dependency to rent capture. Crucially, Doctorow emphasizes that this transition is not simply the result of managerial greed; rather, it is enabled by the erosion of countervailing forces. Competition is neutralized through acquisition. Regulatory oversight is weakened or captured. Labor solidarity dissipates. Interoperability is constrained by intellectual property law. Once these disciplining mechanisms are removed, platforms face diminished consequences for degradation.

Both frameworks illuminate the contemporary collapse of informational trust. Search engines and marketplaces once operated under an implicit epistemic contract: ranking signaled relevance or quality. As sponsored placement and algorithmic steering become central revenue streams, that contract erodes. Users encounter blended feeds in which the distinction between organic signal and paid influence becomes opaque. Sellers must purchase visibility to maintain reach. Attention is allocated not by epistemic merit but by monetized priority.

The economic models thus successfully describe the surface dynamics of platform decay. They identify monopoly leverage, lifecycle degradation, and regulatory retreat as central causal factors. Yet these accounts remain primarily situated within political economy. They explain how extraction occurs, and why it has intensified, but they do not fully articulate why unconstrained optimization reliably converges toward exploitative equilibria across domains.

To move beyond description toward reconstruction requires a more general structural account of optimization, asymmetry, and constraint. The next section develops

such an account by interpreting surveillance and platform optimization through a thermodynamic lens.

3 Surveillance, Gradients, and Thermodynamic Exploitation

Surveillance is not merely a privacy concern; it is a structural transformation of uncertainty into leverage. When behavioral data is collected at sufficient resolution, uncertainty about preferences, vulnerabilities, and constraints is reduced. This reduction of uncertainty generates informational gradients. Optimization engines then exploit these gradients to direct attention, modulate pricing, or adjust compensation.

In thermodynamic terms, a gradient is a structured difference that can be harvested to perform work. Temperature gradients power engines; pressure gradients drive flows. Informational gradients similarly enable economic work. A platform that knows which user is most vulnerable to a given stimulus can target that user with greater precision. A marketplace that knows which seller is most dependent can increase fees incrementally without triggering exit. An employer that knows which worker carries the highest debt burden can adjust wages downward while maintaining labor supply.

Surveillance lowers entropy within the platform’s internal model of the world while increasing entropy for participants. The platform gains predictive clarity; users and workers encounter opacity. Asymmetry in knowledge becomes asymmetry in bargaining power. Dynamic pricing, algorithmic wage discrimination, attention capture, and behavioral nudging are distinct instantiations of the same structural phenomenon: gradient exploitation under weak constraint.

Algorithmic wage discrimination provides a particularly stark example. When staffing platforms purchase credit histories or behavioral data, they are not simply optimizing scheduling efficiency; they are mapping desperation gradients. Workers with higher debt burdens or fewer outside options can be offered lower wages because their acceptance threshold is known. Historically, such fine-grained discrimination would have required prohibitive monitoring costs. Automation renders it scalable. The system thus converges toward individualized extraction calibrated to vulnerability.

Attention markets operate through analogous mechanisms. Recommendation systems collect behavioral signals to model susceptibility to engagement loops. The more precisely susceptibility can be predicted, the more effectively content can be sequenced to sustain attention. Engagement optimization becomes a form of cognitive gradient

harvesting. Each scroll, pause, or click reduces uncertainty, enabling further refinement of the extraction vector.

These processes are not accidental side effects but structurally predictable outcomes of unconstrained optimization. When the objective function is profit maximization and the system has access to high-resolution data, optimization converges toward the steepest available gradient. Without boundary conditions limiting exploitation of asymmetry, the attractor state is extractive equilibrium.

This dynamic aligns with a broader structural insight: optimization is directional. It does not merely increase efficiency; it moves systems along specific vectors defined by their objective functions. In physical systems, unconstrained gradient descent leads toward entropy maximization or energy minimization depending on context. In socio-technical systems governed by profit objectives, unconstrained gradient descent leads toward concentration of leverage and externalization of cost.

The informational trust collapse observed in contemporary platforms can thus be interpreted as a thermodynamic inevitability under current objective functions. When ranking position becomes monetizable, epistemic order transforms into a revenue surface. When attention is priced, cognition becomes a yield curve. When labor metrics are continuously monitored, human dignity becomes a variable in a performance function.

The problem is not optimization per se but optimization without invariant constraints. In the absence of boundary conditions preserving autonomy, reversibility, and reciprocity, systems exploit gradients until social cohesion erodes. The next section situates this convergence within a constraint-first ontology, drawing on RSVP and entropy-bounded governance to articulate a deeper structural diagnosis.

4 Constraint Erosion and the Pathology of Unbounded Optimization

The preceding analysis suggests that extraction and lifecycle degradation are not isolated failures of corporate virtue but manifestations of a deeper structural drift. To understand that drift requires a shift in ontological framing. Rather than treating markets as neutral arenas occasionally distorted by monopoly, a constraint-first perspective asks what boundary conditions define the permissible trajectories of optimization within a system.

Constraint-first ontology begins from the premise that structure precedes content. Systems remain healthy not because their participants are benevolent but because invariant constraints prevent destabilizing gradients from dominating. In physical systems, conservation laws and boundary conditions stabilize dynamics. In socio-technical systems, interoperability norms, regulatory oversight, labor bargaining power, and informational transparency serve analogous functions.

The contemporary internet exhibits what may be termed constraint erosion. Over several decades, legal and institutional reforms have weakened or removed many of the mechanisms that once bounded platform behavior. Anti-circumvention law restricts modification and alternative clients. Intellectual property expansion inhibits reverse engineering and protocol-level experimentation. Antitrust enforcement has retreated from structural remedies toward narrow consumer price analysis. Privacy regulation has often defaulted to disclosure regimes that preserve surveillance while redistributing cognitive burden to individuals.

These changes do not simply permit larger firms; they alter the feasible space of optimization. When interoperability is suppressed, exit costs rise. When surveillance remains unbounded, informational gradients intensify. When acquisitions preempt emergent competitors, adaptive alternatives fail to mature. Optimization therefore proceeds along steeper extractive vectors with diminished countervailing resistance.

Within the RSVP framework, this can be interpreted as a loss of entropy-balancing scalar fields relative to directional vector flows. Vector flows correspond to directed optimization pressures, such as profit maximization. Scalar smoothing corresponds to distributed constraints that equalize gradients and prevent runaway concentration. When scalar balancing mechanisms weaken, vector flows dominate, producing concentrated vortices of control.

In social terms, such vortices manifest as platform centralization, epistemic distortion, and labor precarity. The system does not collapse immediately; rather, it enters a metastable regime characterized by high extraction and low resilience. Trust erodes gradually. Margins narrow for participants other than shareholders. Informational order becomes increasingly opaque. The system remains operational but brittle.

Doctorow's emphasis on the removal of "discipline" can be reframed within this ontology as the elimination of balancing constraints. Competition, labor solidarity, regulatory oversight, and interoperability previously acted as gradient limiters. Their erosion does not guarantee exploitation, but it increases the probability that optimization converges toward extractive equilibria.

This perspective also clarifies why disclosure-based governance proves insufficient.

Requiring users to consent to surveillance does not alter the gradient structure; it merely legitimizes it procedurally. Consent cannot meaningfully constrain optimization when exit costs are high and alternatives are absent. Constraint must operate at the architectural level rather than at the level of individual choice.

The pathology of the contemporary internet is therefore best understood as unbounded optimization within weakened constraint fields. Extraction is the economic expression of this structural drift. Informational trust collapse is its epistemic expression. Algorithmic management and dynamic pricing are its labor and consumer expressions. The underlying cause is not technological inevitability but objective misalignment combined with constraint erosion.

Having situated the diagnosis within a thermodynamic and constraint-first ontology, the essay now turns to reconstruction. If extraction is the predictable attractor of unconstrained optimization, then structural remedies must reintroduce invariant boundary conditions that preserve degrees of freedom and stabilize gradients.

5 Invariant Constraints and Interoperable Architecture

If the pathology of contemporary platforms arises from unbounded optimization under weakened constraint, then remediation must proceed at the level of structural boundary conditions. The central design question becomes not how to fine-tune engagement metrics or rebalance advertising markets, but which invariants must be preserved to prevent extractive convergence.

A constraint-first reconstruction begins with interoperability. Protocol-level interoperability lowers switching costs and restores competitive discipline without requiring complete market fragmentation. When alternative clients, feeds, ranking systems, and filtering layers are legally and technically permitted, users regain degrees of freedom. Interoperability transforms a monolithic optimization engine into a contested ecology of operators. It reintroduces friction against unilateral objective functions by allowing participants to route around them.

Interoperability does not eliminate optimization; it pluralizes it. Competing ranking algorithms, privacy-preserving clients, and alternative interface designs generate distributed objective functions within a shared substrate. The resulting system resembles an ecosystem rather than a vertically integrated extraction pipeline. In thermodynamic terms, interoperability diffuses steep gradients by enabling lateral

flows rather than enforcing singular channels.

Reversibility constitutes a second invariant constraint. Systems should permit exit without catastrophic loss of social or economic capital. Data portability, standardized identity protocols, and open communication layers reduce the cost of departure. When exit is feasible, degradation imposes measurable cost on the degrading party. Reversibility thus reinstates a form of structural discipline without presupposing constant regulatory intervention.

Auditability provides a third constraint. Algorithmic systems that mediate information, labor, or pricing must be subject to meaningful inspection. Black-box optimization accelerates extractive drift by obscuring gradient structure. Auditability introduces transparency into the objective function, allowing misalignment to be detected before it becomes systemic. This need not entail full public disclosure of proprietary algorithms; rather, it requires enforceable mechanisms for verifying compliance with defined boundary conditions.

A fourth invariant concerns bounded surveillance. Surveillance fuels gradient harvesting by reducing uncertainty asymmetrically. Privacy reform should therefore be reframed as optimization constraint rather than consumer preference management. Limiting data accumulation increases entropy within the platform’s predictive model, reducing the precision with which vulnerability can be exploited. In thermodynamic language, privacy introduces controlled noise that prevents predatory gradient amplification.

These architectural constraints are mutually reinforcing. Interoperability lowers exit costs. Reversibility ensures degradation has consequences. Auditability reveals objective misalignment. Privacy limits gradient accumulation. Together they establish a structural field within which optimization may occur without converging toward exploitation.

Such constraints need not eliminate profit motives or competition. Rather, they bound the feasible region of objective space. Firms may compete on service quality, efficiency, or innovation, but not on unrestrained surveillance or coercive lock-in. The aim is not to suppress optimization but to ensure that its trajectories remain compatible with autonomy, dignity, and epistemic trust.

The next section extends this reconstruction by applying entropy-bounded governance and operator-ecological principles to labor, attention markets, and democratic oversight.

6 Entropy-Bounded Governance and Operator Ecology

Architectural constraints address the structural conditions of platform optimization. Yet stable reform also requires governance principles capable of bounding extraction across domains such as labor, attention, and marketplace exchange. Entropy-bounded governance provides a conceptual framework for articulating these limits.

Within an entropy-bounded model, informational and behavioral data are treated as civic resources embedded in a shared ecosystem. Surveillance accumulation increases informational asymmetry, which in turn enables gradient exploitation. When asymmetry surpasses socially sustainable thresholds, the resulting extraction degrades trust and agency. Governance must therefore establish upper bounds on asymmetry accumulation rather than merely regulating downstream harms.

In labor markets, this entails limiting vulnerability-based wage discrimination and biometric surveillance. Data collection should be restricted to operationally necessary metrics. Predictive modeling based on personal debt, health status, or other extrinsic vulnerabilities introduces steep gradients that convert private hardship into employer leverage. Entropy-bounded labor governance constrains such gradients by narrowing permissible data inputs and prohibiting compensation models calibrated to desperation.

In attention markets, entropy bounds require limiting manipulative personalization. Recommendation systems may sort and filter content, but they should not be permitted to exploit high-resolution psychological modeling to induce compulsive engagement. Dark patterns, exploitative notification design, and asymmetrical reinforcement loops represent gradient intensifiers that degrade cognitive autonomy. Bounded governance would constrain these intensifiers through design standards and enforceable limits on behavioral targeting.

Marketplace ranking systems similarly require entropy bounds. When visibility can be purchased without clear demarcation, epistemic order collapses into monetized influence. Sponsored placement may remain permissible, but it must be distinctly partitioned from organic ranking signals. Informational clarity restores epistemic integrity by preventing covert blending of revenue-driven ordering with quality-based ordering.

These domain-specific applications reflect a broader operator-ecological perspective. Platforms are not isolated firms but ecological operators within a socio-technical field. Healthy ecosystems exhibit diversity, redundancy, and modularity. Monocultures

are efficient but brittle; they amplify localized perturbations into systemic failures. Operator ecology therefore prioritizes decentralization of control surfaces, encouraging distributed innovation rather than centralized gatekeeping.

Ecological resilience arises when multiple pathways exist for communication, exchange, and labor coordination. Interoperability fosters such diversity. Modular protocols prevent single points of failure. Public digital infrastructure can serve as a baseline substrate upon which private innovation occurs without monopolistic enclosure. In this model, the platform ceases to be the sovereign arbiter of exchange and becomes one operator among many within a shared environment.

Democratic governance functions as gradient stewardship within this ecology. The central political question is not whether optimization should occur but who defines its objective function and what constraints delimit it. Public institutions must articulate boundary conditions that preserve autonomy and informational integrity. This involves legislative clarity, regulatory enforcement, and institutional capacity to adapt as technological gradients evolve.

Critically, entropy-bounded governance avoids the false dichotomy between innovation and regulation. Innovation proceeds within defined constraint fields, much as engineering operates within material and thermodynamic limits. By establishing invariant constraints, governance channels innovation toward constructive trajectories rather than extractive attractors.

The final section returns to the philosophical horizon of the discussion, addressing whether maximal optimization constitutes a desirable social endpoint and clarifying the normative commitments underlying constraint-bounded reconstruction.

7 Conclusion: Against Maximal Optimization

The crisis of the contemporary internet is frequently framed as a story of corporate excess, technological hubris, or regulatory inertia. Each of these elements is present. Yet beneath them lies a more fundamental structural issue: the normalization of maximal optimization without invariant constraint.

Optimization is often treated as synonymous with progress. Markets optimize allocation; algorithms optimize engagement; logistics systems optimize throughput. Efficiency appears self-justifying. However, efficiency is always relative to an objective function. When the objective function is profit under weak constraint, optimization converges toward extraction. When surveillance remains unbounded, informational

gradients intensify. When interoperability is suppressed, exit costs rise. Under such conditions, the attractor state of the system is not mutual flourishing but asymmetric leverage.

The question raised implicitly throughout the public debate is not merely economic but philosophical. Should every second of human attention be monetizable? Should every asymmetry in vulnerability be harvestable? Should ranking systems that mediate epistemic access be governed primarily by advertising yield? A society committed to maximal optimization along these vectors will produce measurable output, but it will also erode trust, autonomy, and democratic capacity.

Constraint-bounded reconstruction offers an alternative trajectory. By restoring interoperability, ensuring reversibility, limiting surveillance, mandating auditability, and bounding asymmetry in labor and marketplace governance, it becomes possible to stabilize gradients and preserve degrees of freedom. Optimization need not be abolished; it must be situated within structural fields that protect agency and epistemic integrity.

The decay of the internet is therefore neither technological destiny nor irreversible decline. It is the consequence of objective misalignment compounded by constraint erosion. Platforms did not inevitably transform into extractive regimes; they did so because boundary conditions weakened while optimization intensified. Reintroducing those constraints is not an act of nostalgia but an act of architectural correction.

Digital infrastructures shape perception, exchange, and labor at civilizational scale. Their governance cannot be reduced to consumer preference or shareholder return. If informational ecosystems are treated as thermodynamic systems requiring balanced gradients, then stewardship replaces exploitation as the organizing principle.

The task before contemporary governance is not to suppress innovation but to define the invariant constraints within which innovation unfolds. Without such constraints, extraction remains the dominant attractor. With them, it becomes possible to cultivate digital systems oriented toward autonomy, resilience, and collective intelligence rather than toward maximal yield.

In this sense, the future of the internet depends less on technological breakthroughs than on institutional design. The choice is not between efficiency and stagnation but between unbounded optimization and constraint-aligned development. Only the latter preserves the structural conditions under which freedom, trust, and democratic deliberation can endure.

Appendices

A Optimization Under Asymmetric Information Gradients

This appendix formalizes the structural claim that unconstrained optimization under informational asymmetry converges toward extractive equilibria.

A.1 A.1 System Setup

Let there be a platform P interacting with a population of agents $i \in \{1, \dots, n\}$.

Each agent possesses:

$$x_i \in \mathcal{X}$$

representing behavioral state (preferences, vulnerabilities, time constraints, etc.).

The platform possesses a predictive model:

$$\hat{x}_i = f(D_i)$$

where D_i is accumulated surveillance data on agent i .

Define the informational asymmetry for agent i as:

$$\Delta_i = H(x_i) - H(x_i \mid D_i),$$

where H denotes Shannon entropy.

Thus Δ_i measures entropy reduction achieved through surveillance.

A.2 A.2 Profit Optimization

Suppose the platform selects an action a_i (price, wage, ranking weight, content sequencing) to maximize profit:

$$\Pi = \sum_{i=1}^n \pi(a_i, x_i),$$

subject to behavioral response constraints.

If P has perfect knowledge of x_i , it can choose:

$$a_i^* = \arg \max_a \pi(a, x_i).$$

Under uncertainty, expected profit becomes:

$$\mathbb{E}[\pi(a_i, x_i)].$$

As $\Delta_i \rightarrow \max$, predictive uncertainty shrinks:

$$\text{Var}(x_i \mid D_i) \rightarrow 0,$$

and the platform approaches first-degree price or wage discrimination.

A.3 A.3 Extractive Attractor

Define the agent surplus for i as:

$$S_i = U_i(x_i) - C_i(a_i),$$

and platform surplus as:

$$\Pi_i = R_i(a_i) - K_i.$$

Under competitive conditions:

$\max \Pi_i$ is constrained by exit.

Under lock-in and high Δ_i , the feasible action set expands:

$$a_i \in \mathcal{A}_{\text{locked}} \supset \mathcal{A}_{\text{competitive}}.$$

The equilibrium shifts toward:

$$a_i^* \rightarrow \arg \max_a \Pi_i \quad \text{subject only to minimal retention constraint.}$$

Thus extractive equilibrium arises when:

$$\Delta_i \uparrow, \quad \text{exit cost} \uparrow, \quad \text{constraint field} \downarrow.$$

A.4 A.4 Constraint-Bounded Optimization

Introduce a constraint functional:

$$\mathcal{C}(a_i, D_i) \leq \kappa,$$

where κ bounds permissible asymmetry exploitation.

Then the optimization becomes:

$$\max_{a_i} \sum_i \pi(a_i, x_i) \quad \text{subject to} \quad \mathcal{C}(a_i, D_i) \leq \kappa.$$

If \mathcal{C} penalizes gradient exploitation proportional to Δ_i , then:

$$\frac{\partial \Pi}{\partial \Delta_i}$$

is bounded.

Hence bounded surveillance limits convergence to extractive equilibria.

A.5 A.5 Structural Interpretation

Extraction is therefore not anomalous but the natural attractor of:

$$\max \Pi \quad \text{with} \quad \Delta \rightarrow \max \quad \text{and weak constraints.}$$

Reform must operate by:

Bounding Δ , Lowering exit costs, Increasing constraint strength.

This formalization demonstrates that gradient harvesting under informational asym-

metry predicts extractive convergence absent invariant constraints.

B Thermodynamic Interpretation of Surveillance and Entropy Flow

This appendix formalizes the claim that surveillance-driven platforms operate as gradient-harvesting systems that convert informational asymmetries into economic work. The analysis draws analogies from statistical thermodynamics while remaining within an information-theoretic framework.

B.1 B.1 Informational Entropy and Uncertainty

Let X denote the state of an agent (preferences, constraints, behavioral tendencies). The platform maintains a probabilistic model $P(X)$.

The Shannon entropy of the agent state is:

$$H(X) = - \sum_x P(x) \log P(x).$$

After surveillance data D is acquired, the conditional entropy becomes:

$$H(X | D).$$

Define entropy reduction:

$$\Delta H = H(X) - H(X | D).$$

This reduction represents the informational gradient accumulated by the platform.

B.2 B.2 Economic Work from Informational Gradients

In thermodynamics, usable work is extracted from gradients (e.g., temperature differences). Analogously, economic work is extracted from informational gradients.

Define platform profit contribution from agent i as:

$$W_i = g(\Delta H_i),$$

where g is monotonically increasing.

Under zero informational asymmetry:

$$\Delta H_i = 0 \Rightarrow W_i = W_{\text{baseline}}.$$

As ΔH_i increases, predictive accuracy correspondingly improves, thereby expanding the feasible precision with which the platform can calibrate interventions directed at agent i . Formally, this relationship may be expressed as

$$\frac{\partial W_i}{\partial \Delta H_i} > 0,$$

indicating that the marginal extractable work or revenue associated with agent i is strictly increasing in the reduction of informational entropy. Greater informational resolution permits the platform to approximate individualized thresholds with increasing fidelity. In pricing contexts, this manifests as dynamic adjustment toward an agent's willingness-to-pay boundary. In labor markets, it appears as wage calibration sensitive to measured vulnerability or reservation wage proxies. Within attention markets, enhanced predictive granularity enables sequencing strategies that maximize engagement by targeting susceptibility gradients. In each case, informational entropy reduction is converted into economic work through increasingly fine-grained alignment between predicted behavioral response and optimization objective.

B.3 B.3 Entropy Export and Externalization

Let total system entropy be:

$$S_{\text{total}} = S_{\text{platform}} + S_{\text{agents}}.$$

Surveillance reduces internal entropy of the platform model:

$$S_{\text{platform}} \downarrow.$$

However, this reduction of entropy within the platform's internal predictive model frequently corresponds to an increase in effective behavioral entropy at the level of

agents. Formally, one may express this asymmetry as

$$S_{\text{agents}} \uparrow,$$

where S_{agents} denotes the dispersion, instability, or unpredictability imposed upon agents' lived conditions. This increase may take multiple structurally analogous forms. In cognitive domains, continuous engagement optimization can induce attentional fragmentation and overload, generating unstable decision environments. In labor markets, vulnerability-sensitive wage calibration can produce income volatility and economic precarity. In informational ecosystems, monetized ranking can degrade epistemic coherence by displacing relevance-based ordering with revenue-weighted visibility.

The result is a transfer rather than an elimination of entropy. Informational order becomes concentrated within the platform's predictive apparatus, while experiential disorder is distributed outward across the agent population. Predictive certainty is centralized; uncertainty is externalized.

B.4 Gradient Instability and Extractive Vortices

Let G denote an asymmetry gradient (knowledge, bargaining power, attention leverage).

Define system stability condition:

$$|\nabla G| \leq \gamma_{\text{crit}}.$$

If:

$$|\nabla G| > \gamma_{\text{crit}},$$

feedback amplification occurs, producing a concentration vortex in which:

$$\frac{dG}{dt} > 0 \quad \text{and} \quad \text{variance in surplus distribution increases.}$$

Empirically, this dynamic manifests in observable structural shifts across economic and informational domains. Profit margins derived from intermediation tend to increase as platforms refine their capacity to appropriate surplus from dependent participants. Informational trust correspondingly declines as ranking systems and recommendation engines become visibly subordinated to monetized placement rather than epistemic

relevance. Simultaneously, control over visibility, distribution, and access becomes increasingly concentrated within a small number of centralized operators, amplifying their capacity to modulate flows of attention, labor, and exchange. These outcomes are not isolated symptoms but coordinated expressions of gradient amplification under weakened constraint.

B.5 B.5 Entropy-Bounded Governance

Introduce an entropy bound:

$$\Delta H_i \leq \eta,$$

where η defines the maximum permissible asymmetry.

Alternatively, impose gradient damping:

$$\frac{dG}{dt} = -\lambda G,$$

with $\lambda > 0$ representing regulatory or architectural constraint strength.

Under sufficient damping:

$$G \rightarrow G_{\text{stable}},$$

preventing runaway extraction.

B.6 B.6 Structural Implication

Surveillance accumulation is thermodynamically analogous to increasing free energy gradients. Without damping mechanisms (privacy, interoperability, auditability), these gradients drive work extraction until system brittleness increases.

Healthy informational ecosystems require bounded asymmetry and distributed entropy flow rather than concentrated gradient harvesting.

This thermodynamic framing clarifies why privacy reform and interoperability are not merely ethical preferences but structural stabilizers that prevent extractive gradient

amplification.

C RSVP Scalar–Vector Interpretation of Platform Optimization

This appendix situates platform extraction within the Relativistic Scalar–Vector Plenum (RSVP) framework, interpreting socio-technical optimization as a coupled scalar–vector field dynamic.

C.1 C.1 Field Representation

Within RSVP, system dynamics are modeled as interactions between:

$$\Phi(x, t) \quad (\text{scalar field})$$

representing distributed structural density (norms, trust, reciprocity, epistemic integrity),

and

$$\vec{v}(x, t) \quad (\text{vector field})$$

representing directed flows (optimization pressures, profit gradients, attention steering).

Let entropy density be:

$$S(x, t).$$

Healthy systems maintain coupling between scalar smoothing and vector flow.

C.2 C.2 Unbounded Vector Dominance

Platform profit optimization can be modeled as a strong vector field:

$$\vec{v} = -\nabla\Pi,$$

where Π is the profit potential landscape.

In absence of balancing scalar constraints:

$$\frac{\partial \Phi}{\partial t} \ll |\vec{v}|,$$

vector flows dominate system evolution.

The consequences of such vector dominance appear in the reconfiguration of informational structure at scale. Attention flows become increasingly concentrated, channeled through a narrow set of highly amplified nodes rather than distributed across diverse pathways. Distributed epistemic density correspondingly collapses, as heterogeneous sources of relevance are displaced by centrally optimized ranking regimes. Informational topology acquires increasing curvature, with gradients steepening around monetized attractors and flattening elsewhere. What emerges is not merely inequality of visibility but a structural reshaping of the communicative field, in which flow, relevance, and discovery are progressively subordinated to concentrated optimization pressures.

C.3 Scalar Smoothing as Constraint Field

Introduce scalar constraint potential:

$$\mathcal{U}(\Phi) = \alpha \int (\nabla \Phi)^2 dx,$$

where α encodes constraint strength.

Coupled dynamics become:

$$\frac{\partial \Phi}{\partial t} = D \nabla^2 \Phi - \beta \nabla \cdot \vec{v},$$

$$\frac{\partial \vec{v}}{\partial t} = -\gamma \nabla \Phi - \delta \vec{v}.$$

In this formulation, D denotes the diffusion coefficient associated with the scalar field, capturing the rate at which normative structure, trust, and distributed constraints smooth local irregularities across the system. The parameter γ represents the coupling strength between scalar density and vector flow, determining the extent to which directed optimization pressures remain responsive to distributed structural conditions.

The parameter δ encodes vector damping, limiting the persistence and amplification of directed extraction flows.

When γ and δ are small, vector dynamics become weakly coupled to scalar structure and insufficiently damped. Under such conditions, optimization pressures propagate with minimal resistance, and localized gradients intensify into extractive vortices characterized by concentrated control and surplus appropriation. Conversely, when coupling and damping parameters are sufficiently large, scalar balancing forces counteract directional amplification. Diffusion and feedback restrain curvature in the informational field, preventing runaway concentration and stabilizing the system within bounded equilibria.

C.4 C.4 Extractive Vortex Condition

Define curvature threshold:

$$\kappa = |\nabla \cdot \vec{v}| .$$

If:

$$\kappa > \kappa_{\text{crit}},$$

localized collapse of scalar density occurs:

$$\Phi \rightarrow \Phi_{\min}.$$

In socio-technical terms, this structural instability manifests as a collapse of trust within the mediated environment, as participants increasingly perceive ranking, pricing, and allocation mechanisms to be opaque or adversarial. Economic margins for dependent participants compress as platforms appropriate surplus through intensified intermediation and fee extraction. Simultaneously, epistemic quality degrades, as relevance-based ordering is displaced by monetized prioritization and engagement-optimized sequencing. These outcomes are not discrete failures but coordinated expressions of concentrated gradient amplification operating within weakened constraint fields.

C.5 C.5 Interoperability as Boundary Condition

Interoperability can be modeled as a boundary constraint:

$$\Phi(x, t) \geq \Phi_{\min},$$

enforced through:

Open protocol interfaces \Rightarrow Distributed vector fields.

This prevents singular concentration of \vec{v} .

C.6 C.6 Structural Interpretation

Within RSVP terms, platform decay represents vector-field overdominance relative to scalar balancing. Extraction corresponds to curvature amplification in profit gradients absent diffusion.

Reconstruction therefore requires:

$$\gamma \uparrow, \quad \delta \uparrow, \quad D \uparrow.$$

That is, stability requires an increase in scalar coupling, understood as the strengthening of normative and legal constraints that bind directed optimization pressures to distributed structural conditions. It further requires greater vector damping, achieved through mechanisms such as privacy limits and enforceable auditability, which attenuate the amplification of extractive flows. Finally, enhanced diffusion must be introduced through interoperability and modularity, permitting gradients to disperse across multiple pathways rather than concentrating within singular nodes. Together, these adjustments restore balance between directed optimization and distributed structural integrity, preventing runaway concentration and preserving systemic resilience.

This scalar–vector formulation embeds platform extraction within a general field-theoretic model of structural imbalance, clarifying why unconstrained optimization converges toward concentrated control and how boundary conditions restore systemic stability.

D EBSSC Formal Model of Attention and Semantic Extraction

This appendix develops a formal model of attention markets and semantic ordering under the Entropy-Bounded Sparse Semantic Calculus (EBSSC) framework. EBSSC imposes admissibility constraints on semantic ordering operators, bounding extraction by restricting distortion and concentration in exposure distributions. The aim is to formalize how ranking systems become extractive when semantic ordering is subordinated to monetized optimization.

D.1 Semantic Space and Attention Priors

Let (\mathcal{S}, μ) denote a semantic space equipped with a base measure μ (counting or continuous as appropriate). Elements $s \in \mathcal{S}$ represent content items such as posts, products, or media objects.

Each agent i possesses a baseline attentional prior:

$$A_i : \mathcal{S} \rightarrow [0, \infty),$$

normalized with respect to μ :

$$\int_{\mathcal{S}} A_i(s) d\mu(s) = 1.$$

The platform applies a ranking weight function:

$$R : \mathcal{S} \rightarrow [0, \infty), \quad R(s) \geq 0.$$

Exposure probability for agent i becomes:

$$E_i(s) = \frac{R(s)A_i(s)}{\int_{\mathcal{S}} R(s')A_i(s') d\mu(s')}.$$

Thus A_i encodes attentional predisposition, while R reweights exposure through platform ordering.

D.2 D.2 Neutral and Monetized Ranking Operators

Define neutral ranking operator:

$$R_{\text{neutral}}(s) = F(q(s), r(s)),$$

where $q(s)$ denotes epistemic quality and $r(s)$ relevance.

Define monetized ranking operator:

$$R_{\text{monetized}}(s) = G(b(s), y(s), \sigma_i(s)),$$

where $b(s)$ is bid price, $y(s)$ engagement yield, and $\sigma_i(s)$ susceptibility for agent i .

For distortion comparison, normalize ranking weights:

$$\int_{\mathcal{S}} R_{\text{neutral}}(s) d\mu(s) = \int_{\mathcal{S}} R_{\text{monetized}}(s) d\mu(s) = 1.$$

D.3 D.3 Semantic Distortion Metric

Define semantic distortion functional:

$$\mathcal{D} = \int_{\mathcal{S}} |R_{\text{neutral}}(s) - R_{\text{monetized}}(s)| d\mu(s).$$

This is an L^1 distance on normalized ranking weights.

Define epistemic trust as a decreasing function of distortion:

$$T_{\text{epistemic}} = \Phi(\mathcal{D}), \quad \Phi'(\mathcal{D}) < 0.$$

Thus increasing distortion reduces epistemic trust.

D.4 D.4 Aggregate Revenue Functional

Define aggregate platform revenue:

$$\Pi = \sum_i \int_{\mathcal{S}} \rho(s) E_i(s) d\mu(s),$$

where $\rho(s)$ denotes revenue per exposure.

Platform objective:

$$\max_{R(s)} \Pi.$$

Subject to retention constraint:

$$P_{\text{exit}} = \Psi(\mathcal{D}, H_{\text{exposure}}, \text{latency}, \text{friction}), \quad P_{\text{exit}} \leq \epsilon.$$

D.5 D.5 Revenue-Weighted Convergence

In the absence of distortion and diversity constraints, and assuming monotonic correlation between monetization and engagement yield,

$$R(s) \propto \rho(s),$$

leading to revenue-weighted ordering dominance.

D.6 D.6 Exposure Entropy and Concentration

Define exposure entropy (diversity measure):

$$H_{\text{exposure}} = - \int_{\mathcal{S}} E_i(s) \log E_i(s) d\mu(s).$$

High entropy indicates distributed exposure; low entropy indicates concentration.

Define concentration metric:

$$C = \int_{\mathcal{S}} E_i(s)^2 d\mu(s).$$

As entropy decreases, concentration increases.

To distinguish diversity from noise, define relevance-weighted entropy:

$$H_R = - \int_{\mathcal{S}} E_i(s) \log E_i(s) w_q(s) d\mu(s),$$

where $w_q(s)$ weights epistemic quality.

D.7 D.7 EBSSC Admissibility Constraints

Define admissible ranking operator class:

$$R(s) \in \mathcal{R}_{\text{admissible}} \iff \mathcal{D} \leq \theta \quad \text{and} \quad H_{\text{exposure}} \geq H_{\min}.$$

Thus EBSSC constrains ranking operators by bounding distortion and preventing excessive concentration.

Optimization becomes:

$$\max_{R \in \mathcal{R}_{\text{admissible}}} \Pi.$$

D.8 D.8 Attention Capacity and Extractive Intensity

Define collective attention capacity:

$$\mathcal{A}_{\text{total}} = \sum_i T_i,$$

where T_i denotes cognitive time.

Define marginal extractive intensity:

$$\lambda = \frac{d\Pi}{d\mathcal{A}_{\text{total}}},$$

interpreted as marginal revenue per unit collective attention.

Impose civic bound:

$$\lambda \leq \lambda_{\text{civic}},$$

where λ_{civic} denotes the maximum extractive intensity compatible with preservation of public informational goods.

D.9 D.9 Structural Limit Behavior

In the limit of vanishing constraints:

$$\lim_{\theta \rightarrow \infty, H_{\min} \rightarrow 0} H_{\text{exposure}} = H_{\min},$$

and concentration increases toward dominance.

EBSSC does not prohibit monetization but bounds semantic degradation by enforcing operator admissibility.

This refined formalization situates attention extraction within an entropy-bounded semantic calculus, clarifying how ranking systems become extractive and how admissibility constraints stabilize informational ecosystems.

E Gradient-Bounded Labor Surveillance Model

This appendix formalizes labor extraction under surveillance asymmetry and develops a gradient-bounded governance framework consistent with entropy-bounded optimization principles. The objective is to model how informational asymmetry enables wage discrimination and how constraint fields stabilize labor equilibria.

E.1 E.1 Labor State Space and Informational Asymmetry

Let each worker i possess a latent state vector:

$$x_i = (p_i, v_i, c_i),$$

where p_i denotes productivity, v_i vulnerability (e.g., debt burden, outside options),

and c_i constraint profile (time flexibility, mobility, etc.).

Let employer surveillance generate data stream D_i and predictive model:

$$\hat{x}_i = f(D_i).$$

Define informational asymmetry for worker i :

$$\Delta_i = H(x_i) - H(x_i | D_i),$$

where H denotes Shannon entropy.

Higher Δ_i corresponds to greater employer predictive precision.

E.2 E.2 Wage Function Under Surveillance

Let wage offer be:

$$w_i = W(\hat{x}_i).$$

In absence of vulnerability exploitation:

$$w_i = w(p_i),$$

i.e., wages depend only on productivity.

Under surveillance-enabled discrimination:

$$w_i = w(p_i, v_i),$$

with:

$$\frac{\partial w_i}{\partial v_i} < 0.$$

Thus higher vulnerability reduces wage offer.

E.3 E.3 Employer Profit Functional

Employer profit from worker i :

$$\Pi_i = R(p_i) - w_i,$$

where $R(p_i)$ denotes revenue from productivity.

Aggregate employer objective:

$$\max_{w_i} \sum_i \Pi_i.$$

Under high Δ_i , employer approximates:

$$w_i^* = \min\{w : \text{acceptance probability} \geq \alpha\}.$$

Acceptance probability depends on vulnerability:

$$P_{\text{accept}} = \Gamma(w_i, v_i).$$

As $\Delta_i \rightarrow \max$, employer can approximate worker reservation wage:

$$w_i^* \rightarrow w_i^{\text{reservation}}.$$

E.4 E.4 Vulnerability Gradient and Extraction

Define vulnerability gradient:

$$G_v = \frac{\partial w_i}{\partial v_i}.$$

Extraction intensifies as:

$$|G_v| \uparrow.$$

Define surplus distribution:

$$S_i^{\text{worker}} = w_i - c_i,$$

$$S_i^{\text{employer}} = R(p_i) - w_i.$$

As vulnerability discrimination increases:

$$S_i^{\text{worker}} \downarrow, \quad S_i^{\text{employer}} \uparrow.$$

E.5 E.5 Gradient-Bounded Constraint

Introduce asymmetry constraint:

$$\left| \frac{\partial w_i}{\partial v_i} \right| \leq \gamma,$$

where γ bounds permissible vulnerability sensitivity.

Equivalently, restrict wage function domain:

$$w_i = w(p_i),$$

excluding v_i dependence.

Alternatively, impose information constraint:

$$\Delta_i \leq \eta,$$

limiting surveillance accumulation.

E.6 E.6 Stability Condition

Define labor stability functional:

$$\mathcal{L} = \sum_i S_i^{\text{worker}}.$$

Excessive extraction increases variance:

$$\text{Var}(S_i^{\text{worker}}) \uparrow.$$

Instability threshold:

$$\text{Var}(S_i^{\text{worker}}) > \sigma_{\text{crit}} \Rightarrow \text{systemic fragility}.$$

Gradient bounding ensures:

$$\text{Var}(S_i^{\text{worker}}) \leq \sigma_{\text{crit}}.$$

E.7 E.7 Interpretation

Surveillance-enabled wage discrimination represents gradient harvesting along vulnerability axes. Absent constraint, optimization converges toward individualized reservation wages, compressing worker surplus.

Gradient-bounded labor governance stabilizes the system by constraining the channels through which informational asymmetry can be converted into wage discrimination. This may occur through restricting the set of permissible informational inputs available to employers, thereby limiting the accumulation of predictive leverage. It may also involve bounding the sensitivity of wage determination to vulnerability parameters, ensuring that compensation remains tied to productivity rather than to measured precarity. Finally, stabilization can be achieved by strengthening workers' outside options—through collective bargaining, portability, or social safety mechanisms—thereby flattening effective vulnerability gradients v_i and reducing the exploitable differential between reservation wage and offered compensation. In each case, the aim is not to eliminate optimization but to prevent its convergence toward individualized surplus extraction along asymmetry axes.

Thus labor extraction can be modeled as asymmetry amplification under unbounded optimization, and stabilization requires explicit gradient constraints.

This formalization demonstrates that labor surveillance becomes extractive precisely when informational asymmetry enables vulnerability-dependent wage adjustment. Bounding gradients restores structural equilibrium.

F Constraint-Field Stability Theorem

This appendix formalizes the general structural claim underlying the essay: socio-technical systems governed by profit-maximizing optimization converge toward extractive equilibria in the absence of invariant constraint fields. Stability requires bounded gradients and non-vanishing scalar balancing mechanisms.

F.1 F.1 System Dynamics

Let a socio-technical system be represented by state variable:

$$X(t) \in \mathbb{R}^n,$$

encoding distributions of surplus, exposure, informational asymmetry, and control concentration.

Let platform optimization be governed by objective function:

$$\Pi(X),$$

with gradient flow:

$$\frac{dX}{dt} = \nabla \Pi(X) - \nabla \mathcal{C}(X),$$

where $\mathcal{C}(X)$ is a constraint potential.

In the unconstrained case:

$$\mathcal{C}(X) \equiv 0, \quad \frac{dX}{dt} = \nabla \Pi(X).$$

F.2 F.2 Extractive Equilibrium

Define extractive equilibrium X^* as a fixed point satisfying:

$$\nabla \Pi(X^*) = 0,$$

subject to minimal retention constraint:

$$R(X^*) \geq R_{\min},$$

where R measures user or worker retention.

Assume Π is convex along leverage axes (e.g., surveillance accumulation, ranking monetization, vulnerability discrimination). Then gradient ascent increases concentration measures $C(X)$ such that:

$$\frac{dC}{dt} > 0 \quad \text{until constrained by retention.}$$

Thus unconstrained optimization converges to:

$$X_{\text{extractive}}^* = \arg \max_{X \in \mathcal{R}_{\text{retention}}} \Pi(X).$$

F.3 F.3 Constraint-Field Introduction

Let the constraint potential $\mathcal{C} : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ encode invariant structural bounds on extractive amplification by assigning energetic cost to states exhibiting semantic distortion, steep asymmetry gradients, or surplus concentration. Define

$$C(X) = \alpha_1 \mathcal{D}(X) + \alpha_2 G(X) + \alpha_3 V(X), \quad \alpha_k > 0,$$

where α_k are constraint intensities. Assume \mathcal{D} , G , and V are nonnegative functionals on state space, continuously differentiable on the interior of the feasible region, and coercive along the corresponding instability directions, so that $\mathcal{C}(X) \rightarrow \infty$ whenever any of $\mathcal{D}(X)$, $G(X)$, or $V(X)$ diverges.

The distortion term $\mathcal{D}(X)$ may be taken as a normed distance between epistemic and monetized ordering operators, for example an L^1 or L^2 distance between normalized ranking weights, thereby yielding a scale-invariant penalty for order deformation.

The gradient term $G(X)$ is an aggregate magnitude of exploitative differentials, such as a norm of vulnerability sensitivity, information asymmetry, or other leverage gradients, for instance $G(X) = |\nabla \Lambda(X)|$ for a leverage potential Λ . The concentration term $V(X)$ penalizes surplus localization and may be instantiated as a variance functional, a quadratic concentration $\int p^2$, or a Gini-type index, depending on the chosen representation of distributional state in X .

Under these assumptions, $\nabla \mathcal{C}(X) = \alpha_1 \nabla \mathcal{D}(X) + \alpha_2 \nabla G(X) + \alpha_3 \nabla V(X)$ defines a counter-gradient field that introduces curvature opposing extractive ascent directions. Equivalently, \mathcal{C} deforms the effective landscape of the system by raising the energetic cost of trajectories that increase distortion, asymmetry, or concentration, thereby restricting the admissible optimization flow to a bounded region determined by the relative magnitudes of the coefficients α_k .

System dynamics become:

$$\frac{dX}{dt} = \nabla \Pi(X) - \sum_k \alpha_k \nabla \Phi_k(X).$$

F.4 F.4 Stability Condition

A stable equilibrium X^\dagger satisfies:

$$\nabla \Pi(X^\dagger) = \nabla \mathcal{C}(X^\dagger).$$

Linearize dynamics near equilibrium:

$$\frac{d}{dt} \delta X = H_\Pi \delta X - H_C \delta X,$$

where H_Π and H_C are Hessians.

Stability requires eigenvalues of:

$$H_\Pi - H_C$$

to have non-positive real parts.

Thus sufficient stability condition:

$H_C \succeq H_\Pi$ along extractive directions.

F.5 F.5 Extractive Attractor Theorem

Theorem (Extractive Convergence Under Weak Constraint).

Let $\Pi(X)$ be strictly increasing along leverage axes and let $\mathcal{C}(X)$ be negligible or absent. Then for almost all initial states $X(0)$ in the feasible retention region,

$$X(t) \rightarrow X_{\text{extractive}}^*.$$

Proof Sketch. Since $\nabla\Pi$ is non-zero along leverage axes and no opposing gradient exists, trajectories follow steepest ascent until bounded by minimal retention constraint. Convexity ensures convergence to boundary maximizer.

F.6 F.6 Constraint-Bounded Stability Corollary

Corollary. If constraint Hessian dominates profit Hessian along leverage axes, i.e.,

$$\lambda_{\min}(H_C) \geq \lambda_{\max}(H_\Pi),$$

then extractive escalation is prevented and system converges to bounded equilibrium X^\dagger preserving diversity and surplus distribution constraints.

F.7 F.7 Interpretation

The theorem formalizes a general structural principle:

Unbounded profit-gradient dynamics converge toward extractive equilibria constrained only by minimal retention. Stability requires explicit counter-gradients proportional to distortion, asymmetry, and concentration.

In institutional terms, the abstract constraint field $\mathcal{C}(X)$ corresponds to concrete regulatory and architectural mechanisms that introduce curvature opposing extractive ascent directions. Interoperability mandates reduce topological concentration by expanding the feasible adjacency structure of the network, thereby lowering gradient

amplification along centralized nodes. Privacy limits bound informational asymmetry by constraining the admissible domain of surveillance inputs, effectively placing an upper bound on Δ and its induced leverage gradients. Antitrust enforcement alters feasible concentration states by restricting merger-induced increases in \mathcal{C}_{sys} and preserving competitive diffusion. Labor protections constrain the admissible dependence of compensation on vulnerability parameters, flattening wage gradients with respect to asymmetry axes. Auditability requirements introduce observability into the objective function itself, increasing the effective cost of distortion and covert ranking manipulation.

Absent these institutional counter-gradients, the effective constraint potential approaches zero curvature along extractive directions. In such a regime, the system’s trajectory is governed predominantly by the profit-gradient term, and convergence toward extractive boundary equilibria becomes a structural consequence of the dynamical equations rather than a contingent moral deviation. Extraction under weak constraint is therefore not an ethical anomaly but the predictable outcome of unopposed gradient ascent in a high-leverage state space.

This result unifies extraction, surveillance asymmetry, and semantic distortion within a single gradient-flow framework, demonstrating that the pathology of the contemporary internet arises from insufficient constraint curvature relative to optimization pressure.

G Operator Ecology as a Network Dynamical System

This appendix models platforms as operators within a socio-technical network and formalizes ecological resilience in terms of modularity, redundancy, and gradient dispersion. The objective is to show that centralized platform dominance corresponds to topological concentration, while interoperability produces distributed stability.

G.1 G.1 Network Representation

Let the socio-technical ecosystem be represented by a directed graph:

$$\mathcal{G} = (V, E),$$

where nodes $v \in V$ represent operators (platforms, clients, public infrastructure, labor pools), and edges $e \in E$ represent interaction flows (information, capital, labor, attention).

Define adjacency matrix:

$$A \in \mathbb{R}^{|V| \times |V|}.$$

Let flow intensity along edge (i, j) be:

$$F_{ij}(t).$$

G.2 G.2 Centralization and Control Concentration

Define node centrality measure:

$$C_i = \sum_j F_{ij} + \sum_j F_{ji}.$$

High C_i indicates concentration of flows.

Define systemic concentration index:

$$\mathcal{C}_{\text{sys}} = \max_i C_i.$$

When:

$$\mathcal{C}_{\text{sys}} \rightarrow \mathcal{C}_{\text{max}},$$

system approaches monocentric topology.

G.3 G.3 Gradient Amplification in Monocentric Networks

Let profit gradient vector at node i be:

$$\nabla \Pi_i.$$

If a single node dominates flow topology:

$$|\nabla \Pi_i| \gg |\nabla \Pi_j| \quad \forall j \neq i.$$

Thus optimization pressure concentrates at central node.

Define gradient amplification factor:

$$\Gamma = \frac{\max_i |\nabla \Pi_i|}{\text{mean}_j |\nabla \Pi_j|}.$$

As $\Gamma \uparrow$, extractive capacity increases.

G.4 G.4 Interoperability as Edge Augmentation

Introduce interoperability edges:

$$E' = E \cup E_{\text{interop}},$$

adding alternative pathways between nodes.

Define modified adjacency:

$$A' = A + A_{\text{interop}}.$$

Effect on centralization:

$$\mathcal{C}'_{\text{sys}} < \mathcal{C}_{\text{sys}}.$$

Thus interoperability reduces flow concentration and gradient amplification.

G.5 G.5 Resilience and Redundancy

Define resilience functional:

$$\mathcal{R} = \frac{1}{\lambda_{\max}(L)},$$

where L is graph Laplacian.

High $\lambda_{\max}(L)$ indicates tight centralization and vulnerability.

Adding distributed edges lowers $\lambda_{\max}(L)$, increasing resilience.

G.6 G.6 Ecological Stability Condition

System stable if:

$$\Gamma \leq \Gamma_{\text{crit}} \quad \text{and} \quad \mathcal{C}_{\text{sys}} \leq \mathcal{C}_{\text{crit}}.$$

Unstable extractive regime if:

$$\Gamma > \Gamma_{\text{crit}} \quad \text{and} \quad \mathcal{C}_{\text{sys}} \rightarrow \max.$$

G.7 G.7 Interpretation

Centralized platform architectures amplify optimization gradients by concentrating flows through a small subset of nodes. Interoperability functions as edge augmentation, redistributing flows and dampening gradient amplification.

Operator ecology thus furnishes a topological interpretation of constraint-bounded governance. When system architecture is modular, flows are partitioned across semi-autonomous substructures, reducing the maximal centrality of any single node and thereby limiting systemic concentration. When redundancy is present, multiple parallel pathways exist for communication and exchange, increasing robustness under perturbation and lowering the amplification factor of localized gradients. When protocols are open and interoperable, directed optimization pressures are distributed across competing operators rather than channeled through a single chokepoint, diffusing gradient intensity across the network.

Extraction, on this view, is not merely an economic phenomenon but a topological one. It arises when network structure concentrates flows through nodes whose local gradient ascent is insufficiently opposed by alternative routes. Where countervailing pathways are absent or suppressed, optimization pressures accumulate geometrically, producing curvature in the flow topology that reinforces central dominance. Constraint-bounded governance therefore operates by reshaping the network itself, redistributing connectivity so that gradient amplification cannot escalate without encountering structural resistance.

This network formulation complements scalar–vector and entropy-bounded models by demonstrating that institutional architecture determines whether optimization pressures concentrate or diffuse across the ecosystem.

H Unified Gradient-Convergence Theorem for Extractive Systems

This appendix synthesizes the preceding models into a single convergence statement linking informational asymmetry, thermodynamic gradient harvesting, scalar–vector imbalance, and network concentration. The result formalizes extraction as the generic attractor of unconstrained optimization across informational, economic, and topological dimensions.

H.1 H.1 Unified State Representation

Let the socio-technical system be described by the state vector

$$X = (\Delta, G, \Phi, \mathcal{C}_{\text{sys}}),$$

where each coordinate represents a macroscopic structural variable. The quantity Δ denotes aggregate informational asymmetry, understood as the system-level entropy reduction available to dominant operators through surveillance and predictive modeling. The variable G denotes the effective exploitation gradient magnitude, capturing the steepness of leverage differentials through which informational advantage is converted into extractive control over pricing, wages, ranking, or attention allocation. The scalar Φ denotes balancing density, representing the strength of distributed stabilizing structure, including trust, reciprocity, normative constraint, and enforceable institutional

counter-gradients. Finally, \mathcal{C}_{sys} denotes network concentration, measuring the degree to which socio-technical flows are centralized through a limited subset of operators, thereby amplifying optimization pressures via topological chokepoints.

Let platform objective function be:

$$\Pi = \Pi(\Delta, G, \mathcal{C}_{\text{sys}}).$$

Assume:

$$\frac{\partial \Pi}{\partial \Delta} > 0, \quad \frac{\partial \Pi}{\partial G} > 0, \quad \frac{\partial \Pi}{\partial \mathcal{C}_{\text{sys}}} > 0.$$

Thus profit increases with asymmetry, gradient strength, and concentration.

H.2 H.2 Coupled Dynamics

Let system dynamics be governed by:

$$\frac{d\Delta}{dt} = f_1(\Pi) - \alpha_1 \Phi,$$

$$\frac{dG}{dt} = f_2(\Delta) - \alpha_2 \Phi,$$

$$\frac{d\mathcal{C}_{\text{sys}}}{dt} = f_3(G) - \alpha_3 \Phi,$$

$$\frac{d\Phi}{dt} = D \nabla^2 \Phi - \beta G.$$

In this system, the functions f_k are assumed to be monotone increasing, ensuring that amplification along asymmetry or concentration axes reinforces further growth in those directions. The parameters $\alpha_k > 0$ represent constraint coupling strengths, determining the magnitude of countervailing curvature introduced by the constraint field. The coefficient D denotes scalar diffusion, governing the rate at which balancing density Φ spreads across the state space and counteracts localized gradient intensification. The parameter $\beta > 0$ represents the rate at which scalar density degrades under sustained gradient stress, capturing the erosion of trust, reciprocity, or institutional constraint

in the presence of persistent extractive amplification.

H.3 H.3 Extractive Convergence Theorem

Theorem (Unified Extractive Attractor).

If $\alpha_k \rightarrow 0$ and $D \rightarrow 0$ (weak constraint and low scalar diffusion), then for almost all initial states:

$$\Delta(t) \rightarrow \Delta_{\max}, \quad G(t) \rightarrow G_{\max}, \quad \mathcal{C}_{\text{sys}}(t) \rightarrow \mathcal{C}_{\max}, \quad \Phi(t) \rightarrow \Phi_{\min}.$$

Proof Sketch. With α_k negligible, constraint damping vanishes. Since f_k are monotone increasing, positive feedback loops form:

$$\Delta \uparrow \Rightarrow G \uparrow \Rightarrow \mathcal{C}_{\text{sys}} \uparrow \Rightarrow \Pi \uparrow \Rightarrow \Delta \uparrow.$$

Scalar density Φ decreases under sustained gradient stress ($-\beta G$ term). Absent diffusion ($D \rightarrow 0$), Φ cannot recover. Thus trajectories converge to boundary maximizers defined by retention constraints.

H.4 H.4 Stability Under Constraint

If:

$$\alpha_k > 0 \quad \text{and} \quad D > D_{\text{crit}},$$

then equilibrium satisfies:

$$\frac{d\Delta}{dt} = \frac{dG}{dt} = \frac{d\mathcal{C}_{\text{sys}}}{dt} = \frac{d\Phi}{dt} = 0,$$

with bounded values:

$$\Delta \leq \Delta^*, \quad G \leq G^*, \quad \mathcal{C}_{\text{sys}} \leq \mathcal{C}^*, \quad \Phi \geq \Phi^* > 0.$$

Thus constraint coupling and scalar diffusion prevent runaway extractive escalation.

H.5 H.5 Corollary: Institutional Interpretation

The constraint parameters admit direct institutional interpretation. The coefficient α_1 corresponds to privacy and surveillance limits, which bound informational asymmetry and restrict entropy reduction available for exploitation. The parameter α_2 represents labor and asymmetry protections, constraining the admissible dependence of compensation and allocation on vulnerability gradients. The coefficient α_3 encodes antitrust enforcement and interoperability mandates, which limit network concentration and reduce topological gradient amplification. The diffusion parameter D represents normative propagation through democratic governance, institutional learning, and regulatory feedback, enabling scalar balancing density Φ to replenish under stress.

When these parameters approach zero, constraint curvature vanishes along extractive directions, and the coupled dynamical system converges toward boundary maximizers characterized by high asymmetry, steep gradients, and concentrated control. Conversely, when constraint coupling and diffusion parameters exceed critical thresholds relative to amplification rates, bounded equilibria emerge in which asymmetry, gradient magnitude, and concentration remain finite, preserving structural balance within the socio-technical system.

H.6 H.6 Structural Implication

Extraction is not an anomaly but a stable attractor of coupled gradient amplification across informational, economic, and network domains. Stability requires:

$$\text{Constraint coupling strength} \geq \text{Gradient amplification rate.}$$

The health of digital ecosystems is therefore determined not by innovation velocity but by the relative curvature of constraint fields.

This unified theorem demonstrates that the decay of the contemporary internet can be modeled as a dynamical system in which asymmetry accumulation, gradient amplification, and network concentration mutually reinforce under weak constraint. Restoration requires reintroducing invariant structural couplings sufficient to dampen extractive feedback loops.

References

- [1] Arrow, Kenneth J. Economic Welfare and the Allocation of Resources for Invention. In *The Rate and Direction of Inventive Activity*, Princeton University Press, 1962.
- [2] Barabási, Albert-László and Réka Albert. Emergence of Scaling in Random Networks. *Science*, 286(5439):509–512, 1999.
- [3] boyd, danah. *It's Complicated: The Social Lives of Networked Teens*. Yale University Press, 2014.
- [4] Brandeis, Louis D. *Other People's Money and How the Bankers Use It*. Frederick A. Stokes Company, 1914.
- [5] Coase, Ronald H. The Problem of Social Cost. *Journal of Law and Economics*, 3:1–44, 1960.
- [6] Doctorow, Cory. *The Internet Con: How to Seize the Means of Computation*. Verso, 2023.
- [7] Foucault, Michel. *Discipline and Punish: The Birth of the Prison*. Pantheon Books, 1977.
- [8] Gabaix, Xavier and David Laibson. Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *Quarterly Journal of Economics*, 121(2):505–540, 2006.
- [9] Grossman, Sanford J. and Joseph E. Stiglitz. On the Impossibility of Informationally Efficient Markets. *American Economic Review*, 70(3):393–408, 1980.
- [10] Hirschman, Albert O. *Exit, Voice, and Loyalty*. Harvard University Press, 1970.
- [11] Kleinberg, Jon M. Authoritative Sources in a Hyperlinked Environment. *Journal of the ACM*, 46(5):604–632, 1999.
- [12] Landauer, Rolf. Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.
- [13] Lessig, Lawrence. *Code and Other Laws of Cyberspace*. Basic Books, 1999.
- [14] Nadler, Anthony, Matthew Crain, and Joan Donovan. Weaponizing the Digital Influence Machine. *Data Society Research Institute*, 2018.
- [15] Ostrom, Elinor. *Governing the Commons*. Cambridge University Press, 1990.

- [16] Pasquale, Frank. *The Black Box Society*. Harvard University Press, 2015.
- [17] Polanyi, Karl. *The Great Transformation*. Beacon Press, 1944.
- [18] Schumpeter, Joseph A. *Capitalism, Socialism, and Democracy*. Harper & Brothers, 1942.
- [19] Shannon, Claude E. A Mathematical Theory of Communication. *Bell System Technical Journal*, 27:379–423, 1948.
- [20] Stigler, George J. The Economics of Information. *Journal of Political Economy*, 69(3):213–225, 1961.
- [21] Tainter, Joseph A. *The Collapse of Complex Societies*. Cambridge University Press, 1988.
- [22] Varian, Hal R. Computer Mediated Transactions. *American Economic Review*, 100(2):1–10, 2010.
- [23] Wu, Tim. *The Curse of Bigness*. Columbia Global Reports, 2018.
- [24] Zuboff, Shoshana. *The Age of Surveillance Capitalism*. PublicAffairs, 2019.