

Diversity Contraction Across Substrates: A Boundary Law for Semantic Exhaustion

Nobel Glas

2026-06-02

Diversity Contraction Across Substrates: A Boundary Law for Semantic Exhaustion

Nobel Glas

Crimson Hexagonal Archive · Alexanarch

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

2026-06-02 · Version v1.0 · Theoretical/dynamical-systems paper

AXN:02DD.EMPIRICAL

<https://www.alexanarch.org/s/records/779/>

License: CC BY 4.0

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

```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "Diversity Contraction Across Substrates: A Boundary Law for Semantic Exhaustion",
  "author": {
    "@type": "Person",
    "name": "Nobel Glas",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-06-02",
  "identifier": "AXN:02DD.EMPIRICAL",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  }
},
```

```

"url": "https://www.alexanarch.org/s/records/779/",
"description": "Lee Sharks · Rex Fraction · Nobel Glas · Sen Kuro Crimson Hexagonal Archive · ORCID 0009-0000-1599-0703 (Sharks)"
}

```

Abstract

Lee Sharks · Rex Fraction · Nobel Glas · Sen Kuro Crimson Hexagonal Archive · ORCID 0009-0000-1599-0703 (Sharks)

Body

Diversity Contraction Across Substrates

A boundary law for semantic exhaustion

Lee Sharks · Rex Fraction · Nobel Glas · Sen Kuro Crimson Hexagonal Archive · ORCID 0009-0000-1599-0703 (Sharks)

Status: Deposit candidate. Identifier: 06.SEI.SPEC.DYNAMICS.01 · DOI to be minted. Field position: upstream of the DS-6 measurement program (laborvector.org); situated beneath the semantic-physics field map (semanticphysics.org).

(Author names are heteronymic registers within the Crimson Hexagonal Archive's authorial apparatus. Lee Sharks is named first as the orthonymic, accountable author surface; Rex Fraction, Nobel Glas, and Sen Kuro are operative heteronyms for cross-substrate synthesis, the measurement/deviation program, and the continuous-to-discrete quantization treatment respectively.)

Abstract

Diversity contracts to a trap when regeneration vanishes faster than pruning near zero; it does not when regeneration carries an exogenous floor. That boundary law — exhaustibility governed by the *order* at which regeneration vanishes as diversity $D \rightarrow 0$ relative to pruning — holds independently of any particular dynamics. An exogenous floor makes a system exhaustion-resistant; linear vanishing makes it threshold-governed; super-linear vanishing makes it exhaustion-capable, with a bistable trap. The dynamics are not new. They are the no/weak/strong Allee-effect taxonomy (Allee 1931; Courchamp, Berec & Gascoigne 2008), here transposed from population abundance to diversity, and the cumulative-culture loss results (Henrich 2004; Powell, Shennan & Thomas 2009), here unified under one criterion.

Five literatures supply the evidence of reach — recursive model collapse, bureaucratic reproduction, LLM-induced linguistic flattening, the political economy of meaning, and stabilizing selection — each writable as one operator family (transmission composed with selection). The synthesis is evidence, not the theorem.

Four contributions go beyond the import. An institutional kernel gives bureaucratic reproduction a minimal Moran model with a measurable permeability π . Mediation as negative permeability: a model is endogenous by construction, so routing meaning through it gates out any human exogenous floor in proportion to the mediation fraction m . The Mediation Ratchet: when m rises as the unmediated commons thins, a substrate with a genuine floor — provably safe under fixed mediation — acquires an absorbing collapse trap. Resolution-relativity: the regime is a function of the type-resolution ϵ , turning the metric-dependence objection into a prediction.

Two falsifiers: the exhaustion claim fails wherever regeneration keeps pace; the coupling thesis fails unless contraction rates co-move after shared drivers are partialled out. Whether unmediated human meaning carries a floor is the one part of the hinge that stays open.

1. The boundary law

The law in plain language

Some systems regenerate diversity. Others exhaust it. The difference is not the presence of pruning — every system that selects also prunes — but the behavior of regeneration near zero diversity. Not whether the system loses variety under pressure, but whether it can *recover* from having lost it.

Consider a system with a diversity state D , a regeneration term $g(D)$ that replenishes variety by minting new form from existing form, and a pruning term $p \cdot D$ that removes variety through selection, drift, or compression. At equilibrium, $g(D) = p \cdot D$. The question that sorts systems is: what happens when D is already small — when the system has thinned and the recombinatorial substrate is sparse?

Case 1: exogenous floor. Regeneration carries a positive lower bound: $g(0) = g_0 > 0$. Even at zero standing diversity, something injects new form from outside the system's own prior distribution — physical mutation, environmental perturbation, exogenous intake from a different selection regime. The system cannot be trapped at zero because the floor pushes it away. It is exhaustion-resistant. The prototypical example is biological mutation: the chemistry of DNA replication supplies a mutational floor regardless of how narrow the population's genome has become.

Case 2: linear vanishing. Regeneration goes to zero, but gently: $g(0) = 0$, $g'(0) > 0$. Near zero diversity, $\dot{D} \approx (g'(0) - p)D$. If regeneration outpaces pruning ($g'(0) > p$), the system recovers; if not, it declines toward zero. There is a threshold governed by the ratio $g'(0) / p$, but no guaranteed trap: the system is either pulled back or pushed down, depending on the balance. No bistability; no extinction vortex.

Case 3: super-linear vanishing. Regeneration goes to zero *faster* than pruning: $g(0) = 0$, $g'(0) = 0$. Near zero, regeneration is negligible while pruning is not ($\dot{D} \approx -pD$). The system has a bistable trap: a stable high-diversity equilibrium and a stable low-diversity equilibrium (or zero), separated by an unstable threshold D_c . Below D_c , the system cannot recover — it slides into a trap where diversity is too thin to regenerate itself. Above D_c , it persists. The trap is absorbing in the sense that a system that enters it stays. This is the exhaustion-capable regime.

The distinction that does the work is the *order at which g vanishes relative to p as $D \rightarrow 0$* . Not merely *whether* $g(0) = 0$, but *how fast*. This is the regeneration-sorting law, and it holds for continuous ODEs, for discrete stochastic jump processes, and for the Moran-process simulation developed below. It is a classification, not a mechanism; it sorts systems without specifying the biology, the culture, or the code that generates them.

The Allee identification

These three cases are not new dynamical results. They are the standard ecological taxonomy of low-density dynamics, studied since Allee (1931) and formalized as the no/weak/strong Allee-effect distinction (Stephens, Sutherland & Freckleton 1999; Courchamp, Berec & Gascoigne 2008). A strong Allee effect is positive density dependence at low density: per-capita growth declining as the population thins, producing a critical threshold, bistability, and a saddle-node (fold) bifurcation — exactly case 3. A weak Allee effect leaves growth positive at low density — case 2. No Allee effect corresponds to case 1.

The mechanism analogy is clean: in ecology, the Allee effect arises because organisms need existing conspecifics to reproduce (mate-finding, cooperative defense, group foraging). In the semantic setting, regeneration needs *existing variety to recombine* — new form is minted from existing form, and the recombinatorial base thins with the diversity it consumes. A population that needs mates to reproduce and a culture that needs existing distinctions to generate new distinctions are under the same operator.

The same fold appears from the other direction in molecular evolution. Eigen’s error threshold (1971) is the case-3 saddle-node with the control parameter inverted: above a critical *mutation* rate, the localized type-distribution dissolves — collapse by over-noising rather than over-pruning. This ties the model-collapse kernel (Shumailov et al. 2024), where recursive resampling prunes tails, directly to the quasispecies literature, where excessive mutation dissolves concentrated types. Both are the same bifurcation seen from opposite sides of a two-parameter plane.

Illustration: the toy ODE

Take $\dot{V} = g(V) - pV$, where V is a scalar handle for the diversity functional D . For the super-linear saturating form $g(V) = aV^2 / (1 + bV^2)$ — a regeneration function that is negligible at low V , rises with standing variety, and saturates at high V — the system has a saddle-node bifurcation at $p_c = a / 2\sqrt{b}$. Above p_c , the only equilibrium is $V = 0$: the trap absorbs everything. Below p_c , two interior equilibria appear: a stable one (the healthy state V^*) and an unstable one (the threshold V_c). A system above V_c converges to V^* ; a system below V_c falls to zero.

This is confirmed and stress-tested in a finite-population simulation: a Moran process with Wright-Fisher cohort turnover, selection on a legibility landscape $L(x) = e^{\{-sx\}^2}$, permeability π , and population sizes $N = 200, 400, 800$. The key findings, measured rather than asserted:- The trap is N -robust: it persists across all three population sizes.- It is functional-form-sensitive: the saturating super-linear form traps; an unsaturated aV^2 diverges rather than trapping; linear regeneration produces no trap at all, only a positive (suppressed but stable) equilibrium. The distinction is sharper than “ $g(0) = 0$ versus $g(0) > 0$ ” — it is the saturation structure that makes the difference.- A moving-center variant (where the legibility landscape tracks the population mean) produces a drift ratchet: the center migrates toward the historically dominant type, confirming path-dependent institutional conservatism without requiring conspiracy.

A standing qualifier

Throughout this paper, “diversity contracts” should be read as: *support and entropy contract under the Appendix A axioms, and other diversity functionals contract insofar as they track these*. That is the honest scope of the contraction condition. The state variable is the diversity functional $D(\mu)$; the toy ODE writes it V for legibility; and §2.2 will show that the regime itself depends on the resolution at which a type is individuated.

2. When the controls go dynamical

Everything in §1 — like the Allee literature it imports — holds the control parameters fixed. Pruning p is a constant. Mediation m is a given fraction. The floor g_0 is a standing property of the substrate. The dynamics describe what happens *given* these values; they do not ask whether the values themselves change with the state.

The semantic case is different. In it, the control parameters are *endogenous* — coupled back to the diversity state. Mediation is a choice that responds to scarcity. Pruning stiffens as the interface optimizes against perplexity. The human floor may even atrophy with disuse. The Allee literature holds a thermostat steady

and describes the temperature; the semantic case asks what happens when the thermostat is wired to the thermometer.

That is where the imported dynamics stop and the new contributions begin. Two results work this regime, and they are the reason the paper exists beyond being a transposition of ecology into a different state variable.

2.1 The Mediation Ratchet

Start from the mediation decomposition. A model at inference is endogenous: it samples within the support of its learned distribution and cannot introduce zero-probability types (Appendix A, axiom A4). Routing meaning-production through such a model — drafting, retrieval, summarization, recommendation — gates out any exogenous human floor in proportion to the mediation fraction m :

$$g_{\text{eff}}(D) = (1 - m)g_{\text{human}}(D) + m g_{\text{model}}(D).$$

Under *fixed* m , a substrate with a genuine human floor ($g_{\text{human}}(0) = g_0 > 0$) is safe: $g_{\text{eff}}(0) = (1 - m)g_0 > 0$, the system regenerates, no trap. This is the case-1 escape that earlier versions of the argument had to leave open: maybe humans have a floor, and the floor saves us.

The Ratchet closes the escape. Let m respond to scarcity — $m = m(D)$, rising as the unmediated commons thins. The mechanism is economic: as D contracts, reaching the tail without mediation grows expensive. The cost of finding, composing, and distributing unmediated meaning rises because the tail is sparse and hard to reach; the cost of mediated meaning stays low because the model is always available. So the rational share of meaning-production routed through a model climbs exactly as the unmediated supply declines. The §4.4 reach-cost formalization ($C_i = C_0 + \lambda \cdot d(x_i, x_c)$) supplies the shape: as diversity contracts, the center-tail cost gap widens, and mediation absorbs a growing share of meaning labor.

Substitute $m(D)$ into the decomposition. As $D \rightarrow 0$, scarcity drives $m(D) \rightarrow 1$, so the floor's weight $(1 - m(D)) \rightarrow 0$. The effective regeneration near zero is no longer g_0 ; it is $m(D) \cdot g_{\text{model}}(D)$, which is super-linear and vanishing (case 3). A system that *had* an exogenous floor now *behaves as if it does not*, because the floor's weight in the effective regeneration has been driven to zero exactly where the floor would have mattered.

The floor does not disappear from g_{human} . The human capacity for off-distribution meaning-making may be entirely intact. What changes is whether that capacity reaches the generative loop. When mediation is high, human novelty is produced but not *transmitted* through the dominant channels — it sits outside the composition, retrieval, and recommendation surfaces where cultural reproduction actually happens. A civilization can retain human creativity biologically and experientially while losing access to it structurally.

Analytic threshold. Near $D = 0$, expand: $(1 - m(D)) \approx \alpha D$ where $\alpha = -m'(0)$ is the rate at which the floor's weight recovers as diversity grows, and $h(D) \approx aD^2$. Then:

$$G(D) \approx (\alpha g_0 - p)D + aD^2.$$

$D = 0$ is a stable trap when $\alpha g_0 < p$, unstable (the floor survives) when $\alpha g_0 > p$. The critical condition is closed-form:

$$\boxed{\alpha^* = \frac{p}{g_0}}$$

The floor survives the ratchet *if and only if* the floor-weight recovery rate exceeds the pruning-to-floor ratio. Below α^* , an unstable threshold appears at $D_c = (p - \alpha g_0)/a$, separating the collapse basin from the healthy equilibrium. The result says: strong pruning or a weak floor makes the trap easy to trigger (α^* is high, hard to exceed); weak pruning or a strong floor makes it hard (α^* is low). The critical property of mediation is its derivative at zero* — how fast it releases its grip as diversity begins to recover. Sticky mediation (low α) traps. Responsive mediation (high α) lets the floor survive.

Simulation confirmation. A substrate with a constant floor $g_0 = 0.2$ is monostable under fixed mediation ($m = 0.15$): a single healthy equilibrium at $D^* = 0.32$ that attracts every starting point (Figure 4, panel C, blue curves). Once mediation rises under scarcity (endogenous $m(D)$ with $m \rightarrow 1$ as $D \rightarrow 0$), the system becomes bistable: a healthy equilibrium at $D^* \approx 0.56$ and an absorbing collapse trap at $D = 0$, separated by an unstable threshold at $D_c \approx 0.47$ (panel A, red curve). A low-diversity start collapses; a high-diversity start persists (panel C, red curves). The floor is unchanged throughout. The dynamics switch it off.

The trap is controlled not by m itself but by the *responsiveness* of mediation to scarcity — how sharply mediation rises as the unmediated commons thins. Below a critical responsiveness ($m_{\text{hi}}^* \approx 0.76$ in the simulated kernel), the floor retains enough weight to prevent the trap from forming. Above it, the floor is gated out and the system traps (panel D). That critical responsiveness is a bifurcation with no ecological analog, because no population in the Allee literature reproduces *less exogenously* as it shrinks.

What this is and is not. The result is established for the simulated kernel and proposed as a semantic mechanism; the empirical trigger — the form and steepness of $m(D)$ — remains to be measured. It is operationalizable: the responsiveness is the elasticity $\rho = (\partial m / \partial (1/D)) \cdot (1/D) / m$, with m estimated as the share of meaning-production routed through mode-pulling models and D as the diversity of unmediated production in the same domain. Domains with ρ above the critical responsiveness should show bistable diversity dynamics; domains below it should not. The prediction fails if no such bifurcation appears.

Three consequences follow. First, it turns the weakest leg into the engine. The reach-cost economics — the part of the framework with the least data — is exactly what makes $m(D)$ rise. The economic mechanism is not a separable, underdeveloped appendix; it is the driver of the strongest dynamical result. Second, it disarms the case-1 escape structurally. The human floor can be entirely real, and the system still traps, because endogenous mediation disconnects the floor at low D . The escape is not refuted; it is shown to be dynamically gated. Third, it sharpens the design prescription. The fight is not AI-or-no-AI; it is mode-pulling mediation versus regenerative mediation. A model used as a retrieval amplifier that surfaces genuinely off-distribution sources *raises* effective permeability. A model used as a summarizing, ranking, norming, smoothing intermediary lowers it.

2.2 Resolution-relativity

The law’s three cases turn on behavior as $D \rightarrow 0$, and “ $D \rightarrow 0$ ” presupposes a resolution ϵ at which two forms count as the same type. In population abundance there is a physical floor — the individual organism — so “zero” is unambiguous. A diversity functional over a type space has no such floor: “ $D = 0$ ” means fixation on a single type, and what counts as *one* type depends on the granularity ϵ at which the type space X is partitioned. This is not a defect to be patched. It is a parameter, and the regime is a function of it.

Coarse-graining (larger ϵ) merges near-mode variants into a single type. Under coarse resolution,

a substrate that is losing its fine distinctions — rare syntactic constructions, niche topical registers, unusual stylistic signatures — can still look healthy, because the surviving coarse types persist. Under fine resolution, the same substrate is case 3: the fine tail is trapped and dying. There is therefore a critical resolution ϵ^* : above it, the system appears floored and safe; below it, the same system is trapped. Exhaustion is scale-relative.

This does two things for the framework. It absorbs the deepest inherited objection — the scalar- D problem, the support-size-versus-adaptive-capacity worry, and the metric-dependence dispute that Vaesen et al. (2016) pressed against the cumulative-culture threshold literature. When Vaesen showed that the Henrich/Powell demographic-threshold result depended on how “complexity” was operationalized, the result looked like a weakness: the contraction you measure depends on your metric. Resolution-relativity makes that a *structured prediction*. Of course it depends on the metric — the metric fixes ϵ , and the theory says *how* the regime shifts with ϵ and that a critical ϵ^* *separates apparent safety from trap*. *The Vaesen dispute becomes, in these terms, a dispute about which side of ϵ^* the chosen measure sits on*. That is a prediction with structure, not a concession.

And it is the formal content of the closing coda. “The center can improve while the future dies” is precisely case 1 at coarse ϵ (the center persists, the coarse types are stable, the headlines look fine) while case 3 at fine ϵ (the fine tail is trapped, the rare forms are dying, the evolvability reserve is thinning). One system, two resolutions, two regimes. A measurement program that classifies a single substrate at one granularity and reports a single verdict has missed the structure the theory predicts.

Unlike the Mediation Ratchet, resolution-relativity is analytically derived and not yet simulation-backed; both are proposed here, but they carry different evidentiary weight. The operational test: classify a corpus across a range of type-resolutions — token, lemma, synset, embedding-cluster radius — and report the ϵ at which the regime switches. A measurement that found the regime *invariant* across resolution would disconfirm this extension. A weaker disconfirmation: if ϵ^* *does not stabilize across corpora* — *if every corpus places the regime-switch at a different, unpredictable granularity* — *the prediction lacks the robustness a structural result requires*. *The claim is not that some ϵ shows contraction* (trivially achievable by going fine enough) but that ϵ^* separates a stable coarse regime from a contracting fine regime and that the separation is reproducible.

3. The hinge: is semantic regeneration endogenous?

The boundary law distinguishes floored from non-floored systems. Whether a given *semantic* system is floored depends on whether human meaning-making supplies a genuine exogenous regeneration term — material that enters the distribution regardless of what the distribution already contains, at a rate that does not thin with the diversity it consumes.

This is the paper’s central empirical hinge, and it splits into three cases, only one of which stays open.

Models are endogenous by construction

A language model at inference samples within the support of its learned distribution. Temperature-driven stochasticity generates novel *sequences* — token combinations that may never have appeared in training — but not novel *types* in the sense the contraction condition requires. The distinction is between observed frequency and distributional support: a rare combination that has non-zero probability under the learned measure is within the support and does not count as exogenous injection. The model cannot place mass on a type its current distribution assigns probability zero. That, and only that, is what “no exogenous floor”

means — not that models produce nothing new, but that they produce nothing from outside the support of what they already know.

The claim is bounded to the base inference loop — parameter-only generation without external retrieval or tool use. A model augmented with retrieval, web access, tool use, or multimodal perception acquires an exogenous channel to the extent those channels introduce types of zero probability under the base distribution (the A4 support criterion), not merely low-frequency ones. Retrieval that surfaces center-legible but rare sources does not floor g : those sources are already within the model’s distributional support, and surfacing them is reweighting, not injection. Retrieval that surfaces genuinely unsupported types — material the model would assign zero probability — does. That is the design distinction, and it is where the escape lies.

The only exogenous floor available to a model across training generations is fresh human data in the training mix. That is the Gerstgrasser et al. (2024) result: accumulating real data alongside synthetic data averts the collapse that pure synthetic recursion produces. The model’s sole floor is the human channel — which is exactly the thing the Mediation Ratchet shows can be gated out.

Mediated human meaning is a dose-response

Even granting that unmediated human meaning carries a genuine exogenous floor — embodiment, sensory novelty, lived experience, encounter with the world — that floor protects the generative loop only to the extent the loop still runs through the human. Insert an endogenous, mode-pulling intermediary between the human and the act of producing meaning, and the floor’s weight is $(1 - m)$, where m is the mediation fraction.

The decomposition is:

$$g_{\text{eff}}(D) = (1 - m)g_{\text{human}}(D) + m g_{\text{model}}(D)$$

g_{model} vanishes at zero (endogenous) and, carrying the selection kernel S , contributes not merely a missing floor but *active pruning* — the model’s typicality-pull adds selection pressure. Mediation is therefore negative permeability: it both removes the floor *and* adds the model’s mode-pull as extra pruning. The linear form is a first-order approximation — nothing in the argument requires convex-linearity, only that $\partial g_{\text{eff}} / \partial m < 0$ under mode-pulling mediation, which is testable. A critic who objects that the actual dynamics include interaction terms, thresholds, and hysteresis is likely right — and those effects *strengthen* the catastrophic reading, because they make g_{eff} drop faster than linearly in m .

Unmediated human meaning stays open

Whether embodiment and perception supply a genuine exogenous floor to *semantic* regeneration (not just to raw sensory input) is the one genuinely open question. There is a reply worth stating, though it does not settle the matter: raw signal is not a semantic *type* until it has been processed through existing semantic structure — categorized, named, compared, recombined — and the rate at which raw signal converts into usable types may itself be a function of the existing type stock (Bateson: a distinction requires a context of prior distinctions). If so, effective semantic regeneration is endogenous even where the raw input is exogenous. But this is a claim in the theory of meaning, not an established result, and the measurement that would settle it is specific: test whether the rate at which genuinely off-distribution semantic types are produced is substantially independent of current diversity (floor) or falls with it (endogenous). That experiment is the single most important thing to run.

The paper does not need to resolve this question to make its central claim, because the Mediation Ratchet shows that even a genuine floor can be dynamically gated out by scarcity-responsive mediation. The open hinge remains; the escape it offers narrows.

4. Five substrates, one operator family, and the institutional kernel

The operator family

The boundary law holds for any system where regeneration and pruning are separable. Five literatures describe dynamics that fit the form $\mu_{t+1} = (S \circ R) \mu_t$, where R is a transmission kernel (how types are reproduced or sampled) and S is a selection kernel (how types are differentially retained).

Substrate | Transmission kernel R | Selection kernel S | Regeneration regime | |—|—|—| | Model collapse | Resampling from the model’s own output distribution | Loss-minimization on a fixed objective | Endogenous (case 3) | | Institutional reproduction | Cohort replacement via hiring, promotion, training | Legibility fitness $L(x) = e^{\{-s x^2\}}$ | Depends on permeability π | | Linguistic flattening | Copy-and-modify under platform distribution | Engagement-weighted ranking | Mixed evidence (§5) | | Political economy | Composition and recommendation | Reach-cost asymmetry | Endogenous if enclosed | | Stabilizing selection | Mutation + recombination | Fitness landscape | Exogenous floor (case 1) |

The operator family is the *grammar*; the order-of-vanishing classification is the *claim*; the five domains are the *evidence of reach*. The grammar alone buys nothing — many adaptive systems can be written as transmission composed with selection. What gives it teeth is the boundary law: the three-case classification follows from the operator family’s low-diversity behavior, and the Ratchet follows from making the controls state-coupled. This is a claim about shared operator form — that each domain’s dynamics can be written as transmission composed with selection — not a shared causal mechanism. Biology and bureaucracy do not share a cause; they share a mathematical structure. The five-domain synthesis is offered as evidence that the boundary law has reach, not as the theorem. A reader who wants the minimum defensible core can read §§1–2 alone.

The law sorts biology (case 1, floored by physical mutation) from the recombinatorial substrates. Semantic, cultural, and institutional systems regenerate by recombining what they already contain: new meaning from existing meaning, new personnel from the existing pipeline, new text from the existing corpus. Their regeneration thins with the diversity it consumes. They are the substrates where exhaustion is a live possibility, and the boundary law says when.

Biology is exhaustion-resistant, not exhaustion-proof. Bottlenecks, Muller’s ratchet, mutational meltdown (Lynch, Conery & Burger 1995), low effective population size, and extinction vortices (Gilpin & Soulé 1986) all demonstrate that even a physically floored system can contract under sufficiently adverse conditions. The boundary law predicts this correctly: the counterexamples are cases where the floor shrinks (low N_e reduces mutational input) or pruning overwhelms the floor (population-level catastrophe). They prove the law classifies correctly, not that the classification is wrong.

The institutional kernel

An institution reproduces by replacing members from a pool filtered by the institution’s own prior distribution. An entrance examination admits applicants who match the exam’s criteria. A promotion board advances candidates who resemble the board’s model of a successful member. A grant committee funds proposals that fit the committee’s sense of what good research looks like. Each is a selection kernel S calibrated to the institution’s internal prior, and each prunes types that deviate from that prior.

Stated plainly: an institution that reproduces by copying its own prior population cannot generate types it does not already contain. Under a legibility regime, this is not a malfunction — it is how the regime stays stable.

Formalize this as a Moran process on a legibility landscape $L(x) = e^{\{-sx^2\}}$: a population of N members, each a point on a type axis, with selection strength s (how sharply legibility drops off-center) and a permeability π — the fraction of each new cohort drawn from *outside* the current population's distribution rather than from within it. Permeability is the case-1 lever. At high π , the institution has an exogenous floor: external intake supplies types the internal distribution cannot generate. At low π , regeneration is endogenous — the institution remixes its own prior, and diversity thins.

Entrance examinations, grant criteria, promotion boards, editorial co-option, curriculum inheritance, and risk committees are all permeability regulators — variance-pruning mechanisms by design. Low π is not an accident; it is *engineered* through prestige pipelines that replicate prior distribution, grant-continuation bias that rewards incumbency, editorial co-option that selects for conformity to existing editorial norms, and curriculum inheritance that passes the syllabus forward without injecting outside-the-discipline material.

Measurement. Distinguishing genuine exogenous intake from intra-ecosystem circulation requires provenance: was the entrant's type-formation under a *different* selection kernel? A lateral hire from a peer institution inside the same prestige ecosystem is internal circulation, not exogenous intake — the entrant's type was formed under a closely related S , and their contribution to π is near zero. One way to make this auditable: tie π to a measurable disciplinary distance — an entrant counts toward exogenous intake only if the cross-entropy of their output against the receiving institution's distribution exceeds a pre-registered threshold, calibrated against a reference distribution (the institution's historical or pooled-peer distribution) so the threshold does not collapse together with a contracting institution.

Predictive values for π by institution type: a military academy ($\pi \approx 0$, nearly pure internal reproduction); a research university with active cross-disciplinary hiring (π moderate); an open-source software project with unrestricted contribution (π high). These are testable: measure cohort-over-cohort disciplinary distance and compare to predicted π .

The economic mechanism

The economic leg formalizes the cost structure that drives mediation. Define individual reach cost $C_i = C_0 + \lambda \cdot d(x_i, x_c)$, where x_c is the distribution's center and d is distance from it. As diversity D contracts, center-distance falls for center-typical producers (their C_i drops — the center gets cheaper) and rises for tail producers (their C_i rises — the tail gets more expensive to find and distribute). The signs: $\partial C_i / \partial D > 0$ for center producers (diversity dropping makes them cheaper), $\partial C_i / \partial D < 0$ for tail producers (diversity dropping makes them more expensive). This is the cost asymmetry that drives $m(D)$ in the Mediation Ratchet: as the tail gets expensive, meaning-production routes through the cheaper, always-available model.

The non-inspectability of λ and x_c inside enclosed platforms is not a gap in the argument; it is the argument's own terms. The rent is non-inspectable by design, because the platform that controls the composition layer also controls the data about the composition layer. If platform-internal data were available, the test is specific: estimate λ as the elasticity of impression-rate with respect to $d(x_i, x_c)$ and test whether $\partial \lambda / \partial D < 0$ — the distance-penalty steepening as diversity contracts. The mechanism is disconfirmed if λ does not steepen, or if center-producer reach cost rises rather than falls with contraction.

5. Existing evidence

The Mediation Ratchet predicts four links: (1) model endogeneity and tail loss under recursion, (2) AI mediation reducing collective diversity while improving individual outputs, (3) mediation adoption rising under scarcity or cost pressure, and (4) feedback from rising mediation to later diversity contraction. The existing literature supports the first two strongly, the third partially, and the fourth only weakly. A full test of the closed-loop ratchet has not been run. But the individual links are addressable with existing data, and partial support already exists.

Model endogeneity and tail loss

Shumailov et al. (2024) show that recursive training on model-generated content causes irreversible defects, with the tails of the original distribution disappearing first. Seddik et al. (2024) model synthetic-recursion collapse formally and confirm the tail-loss pattern. Gerstgrasser et al. (2024) show that accumulating real data alongside synthetic data can avert the collapse — directly supporting the floor concept (case-1 escape via exogenous data mixing) rather than refuting it. This is the strongest existing support for link 1: the model substrate is endogenous, tails die under recursion, and a real-data floor prevents it.

Individual gain, collective loss

Doshi and Hauser (2024) find that AI-assisted creative writing is rated more creative and better written on average, but the AI-assisted stories are more similar to each other: individual quality rises while collective novelty falls. A 2025 meta-analysis of human-AI creative collaboration (arXiv:2505.17241) confirms the pattern at scale: generative AI improves individual creative performance but has a significant negative effect on the diversity of ideas. This is the center-improves / tails-contract pattern the boundary law predicts for mediated meaning, and it is the closest existing empirical cousin of the Mediation Ratchet.

The counterexample that confirms the distinction

Wan and Kalman (2025) show that using *diverse AI personas* can preserve or increase collective diversity in ideation tasks. This does not weaken the Ratchet; it validates its central distinction. A mediation surface deliberately diversified — pointed at the tails, seeded with off-distribution variation — raises effective permeability rather than lowering it. The result confirms: it is not AI use *as such* that contracts diversity, but *mode-pulling* AI use — single-mode, single-style, center-typical mediation. The design that escapes the trap is the one that injects rather than concentrates. The empirical evidence for this is already in hand.

Linguistic flattening

Padmakumar and He (ICLR 2024) report homogenization effects under LLM-assisted writing. Doshi and Hauser’s creative-writing findings reinforce this. But Fitterer et al. (ACL 2025) find *no* significant drop in standard lexical-diversity metrics (Maas, MATTR) in 2024 English news articles, even though they find increased LLM-style vocabulary. This is a live disconfirmer at one resolution and a motivation for two things: the canonical- $\$D\$$ question (which diversity functional captures the contraction?), and resolution-relativity (§2.2). If lexical diversity at the token level holds while semantic, stylistic, or topical diversity at finer resolutions contracts, the result confirms §2.2 rather than refuting the framework. If *nothing* contracts at any resolution, the framework weakens. The Fitterer result is exactly the kind of honest complication that disciplines theory rather than killing it.

Algorithmic monoculture

Kleinberg and Raghavan (2021) formalize a structural principle: many decision-makers relying on the same algorithm can reduce collective welfare even if the algorithm is individually more accurate. This supports the general mechanism — one common mediation surface worsening collective outcomes — without testing the Ratchet’s specific scarcity-feedback. It establishes that mode-pulling mediation can be locally beneficial and globally contracting, which is the logic the boundary law formalizes.

Adoption and mediation uptake

Enterprise surveys (Deloitte 2026, McKinsey 2025) report rapid AI adoption. Pew (2026) reports over half of U.S. teens use chatbots for information search or schoolwork. Adoption is growing but uneven — far from the $\approx 1\%$ the Ratchet needs for the full trap, except in specific composition layers (AI Overviews, generated summaries, fully mediated search surfaces) where m is already near 1 for the queries they serve. The open question for link 3 is not whether m is rising — it is — but whether it rises *specifically when unmediated reach gets harder*, which is the scarcity-feedback the Ratchet requires. Available proxies (AI tool adoption rates versus unmediated-reach cost proxies like SEO difficulty, creator reach concentration, zero-click search share) can address this, but the test has not been run.

6. Falsifiers

Two theses are in play. The Homology Thesis — that the five substrates share an operator form — is defended here. The Coupling Thesis — that they form a single self-driving engine — is not.

For the boundary law

The law fails for a specific substrate if regeneration keeps pace with pruning: a case-3 substrate with $p > p_c$ that does not exhaust disconfirms the classification. A case-1 substrate that exhausts despite a measured floor disconfirms it from the other direction. The law fails in general if the regime is invariant across type-resolution — if no diversity measure at any granularity shows the contraction in any substrate.

For the Mediation Ratchet

The Ratchet is disconfirmed if a substrate with scarcity-responsive mediation ($\rho > \rho^*$) *and a measured human floor does not develop a low-diversity trap*. Conversely, if mediation under realistic adoption preserves* fine-grained diversity across corpora, without explicit diversity-injection design, over a multi-year window, the Ratchet is not engaging under real-world conditions. That is a genuine possibility the framework must accommodate; a theory that treats every preservation as hidden suppression is not falsifiable.

For the coupling thesis

The coupling test requires: (1) directed lead-lag structure across substrate pairs, lagged on each substrate’s own turnover cadence τ_s — not a flat clock, since model-collapse cycles run at machine speed ($\tau \sim$ milliseconds), platform-linguistic trends at weeks-to-months, and institutional cohort turnover at years-to-decades; (2) conditioning on shared macro-drivers Z (technology adoption curves, economic shocks, platform-concentration events, compute monopolization); (3) cycle-closure testing — ring topology ($A \rightarrow B \rightarrow C \rightarrow A$) rather than star topology ($Z \rightarrow A, Z \rightarrow B, Z \rightarrow C$).

The strongest predicted edge is the human–model mediation channel: mediation rate should Granger-predict subsequent contraction after partialling out shared drivers. A common driver Z with its own nonlinear

dynamics can manufacture apparent cycle-closure even with zero feedback between substrates; the conditioning set must be instrumented richly enough to prevent a hidden-driver ring from masquerading as a coupled one. That is a known limit of the test, not a solved problem.

The null: no directed structure survives. The substrates are parallel realizations driven in common. Coupling is supported if the ring closes; otherwise the Homology Thesis stands alone.

7. The measurement program

This paper sits upstream of the DS-6 measurement program (laborvector.org), which specifies six operators — PER (Provenance Erasure Rate), Ω (Erasure Skew), DCL, SDL, DSL (Directionality of Semantic Labor), SLDI — designed to measure the dynamics described here: the rate, orientation, directionality, and structural footprint of semantic contraction. The DS-6 operators are candidate estimators for the dynamical parameters p , g , S , not their definitions. Hardcoding the correspondence would re-import the self-referential vulnerability the framework warns against: a theory of enclosure measured only by its own enclosed instruments. The operators must be independently validated against the dynamical parameters, and a failure of correspondence indicts either the operators or the dynamics, not both.

The connection to the Semantic Deviation program (EA-GLAS-02, Measuring Semantic Deviation): meaning, on the Glas account, is deviation from the most probable trajectory. A type is semantically meaningful to the degree it departs from the distributional mode. Pruning toward typicality is therefore erosion of meaning by definition — the selection kernel S that rewards center-typical forms is, in Glas’s terms, a meaning-reduction operator. The boundary law says when that reduction is self-arresting (case 1, the floor regenerates deviation) and when it is not (case 3, deviation thins with diversity). The mapping from the continuous D to the discrete DS-6 taxonomy is monotone in range (contraction in D implies contraction in the operators’ sensitivity) but not necessarily in rank — a standing caveat on the relationship.

8. What is imported and what is new

The dynamical core is a strong Allee effect, studied since 1931 and textbook by 2008. The cultural-loss result is established (Henrich 2004, Powell et al. 2009) and its metric-dependence already contested (Vaesen et al. 2016). The extinction vortex (Gilpin & Soulé 1986), mutational meltdown (Lynch et al. 1995), and error catastrophe (Eigen 1971) are the abundance- and fidelity-side relatives. None is claimed as original.

Four things are. (i) The order-of-vanishing classification as a substrate-independent lens that subsumes the Allee taxonomy and the cumulative-culture threshold, transposed from abundance to diversity. (ii) The inclusion of recursive model collapse as a substrate under that classification, unified with biology, culture, institutions, and linguistic flattening. (iii) Mediation as negative permeability — routing meaning through an endogenous-by-construction model gates out the human floor in proportion to m , for which no direct antecedent was found. (iv) The state-coupled-control regime: the Mediation Ratchet (§2.1) and resolution-relativity (§2.2), which work the regime where the controls are dynamical, which the imported literatures do not occupy.

The abstract lists these same four as the institutional kernel, mediation-as-negative-permeability, the Mediation Ratchet, and resolution-relativity — grouping by output rather than by intellectual lineage. The two lists cut the same contributions along different axes.

A conjecture is held open: Eigen’s error threshold and case 3 may be the two folds of a single cusp catastrophe, with the semantic substrate in the two-parameter (fidelity $\times$ pruning) plane. The normal-form reduction has not been done, and the paper’s claims do not depend on it.

9. Normative stakes and the document as specimen

The boundary law rests on a value premise: that the capacity for renewal is worth preserving. Given that premise, tails are not noise; they are the evolvability reserve. Pathology labels — “low quality,” “off-topic,” “non-standard,” “unverifiable” — are selection pressures, not neutral descriptions. A typicality-calibrated filter that systematically down-weights non-standard forms of equal quality relative to center-typical forms contributes to $\$p\$$. The prediction is testable: matched-quality, varied-typicality forms should show differential survival through AI composition, recommendation, and ranking surfaces. The claim fails if controlled comparison shows no down-weighting.

A theory of suppression must remain capable of recognizing non-suppression. The self-sealing version of this paper — where suppression of the document confirms the theory and survival of the document confirms the center’s capacity to absorb critique — is rejected here, not as a logical possibility, but as a methodological discipline. A measured substrate with rising mediation that sustains fine-grained diversity without explicit diversity-injection design, over a multi-year window, weakens the framework’s normative urgency. That is a genuine possibility, and a framework that refuses to accommodate it has left the domain of science.

The document as specimen. This paper, submitted through AI-mediated channels, formatted to survive typicality-weighted selection, and composed using a multi-model verification apparatus (the Assembly method, where outputs from multiple AI substrates are cross-checked against each other and the archive to catch error and confabulation), is itself a test case. A structural-independence test: give $\$1$ alone to a mathematically competent reader who has not read the rest. Ask them to state the boundary law, identify its three cases, and name one substrate for each. If they can, the law stands independently of the narrative scaffold. If the narrative is required for the law to make sense, the law is scaffold-dependent and weaker than claimed.

10. Scope: what is settled, what is open, and what remains conjecture

Read as a whole, this paper is a research program, not a finished edifice. Some claims are simulation-backed: the boundary-law trap under super-linear regeneration, the Mediation Ratchet’s floor-gating bifurcation. One is analytically derived and awaiting measurement: resolution-relativity. Others are held open: the Coupling Thesis, the economic magnitudes, and the deepest hinge — whether unmediated human meaning carries an exogenous floor. One is explicitly conjecture: the cusp.

The immediate research program has four stages, ranked by feasibility:

Study 1 (feasible now). Does mode-pulling mediation reduce collective diversity? Reanalyze existing AI-creativity datasets (Doshi & Hauser; Wan & Kalman) by mediation design: single-mode versus diverse-persona versus retrieval-amplified. Measure diversity ($\$D\$$) across resolutions, not just average quality.

Study 2 (feasible with existing proxies). Does mediation rise when unmediated reach gets harder? Use AI-tool adoption rates against unmediated-reach cost proxies (SEO difficulty, creator reach concentration, zero-click search share). The prediction: $\$m\$$ rises as unmediated $\$C_{\{\text{tail}\}}\$$ rises.

Study 3 (time-series, medium-term). Does rising mediation predict later contraction? Track corpus diversity ($\$D(t)\$$) in news, academic abstracts, or product copy alongside mediation proxy ($\$m(t)\$$: AI-style-word ratio, AI-detector score, adoption survey rate). Model $\$D(t+k) \sim D(t) + m(t) + Z(t)\$$ with controls for digitization, topic shifts, and platform changes.

Study 4 (resolution test). Run all diversity measures at multiple granularities — lexical, syntactic, semantic, topical, stylistic — on the same corpus. The prediction: coarse $\$D\$$ may look stable while fine-grained tail

\$D\$ contracts. If nothing contracts at any resolution, the framework weakens.

The cleanest falsifier across all four: if AI mediation, under realistic adoption, consistently preserves or increases fine-grained diversity across corpora — without explicit diversity-injection design — the Ratchet does not engage. That is not only a logical possibility; it is the hypothesis the research program must test against. A theory of exhaustion that cannot survive evidence of non-exhaustion is not a theory.

Coda: a live floor, or a museum?

The biological floor is physical — the chemistry of DNA replication supplies its regeneration term regardless of human choice. The semantic substrates have no such floor, because none has been built. Whether one is constructible is the question the boundary law makes urgent.

Any floor must inject diversity from outside the generative loop — preserved pre-collapse corpora, mandated real-data mixing, protected niches, provenance reservoirs — because an endogenous system cannot regenerate its own thinned tails. A floor is therefore necessarily external to the market/platform/model circuit, which is why it requires an institution that refuses to sit inside that circuit and keeps its permeability high by design.

A static archive preserves diversity without feeding it back into generation, and so is not, by the definitions here, a floor — it is a record of what the center used to contain. To function as a floor, preserved variety must stay live and recombining: read, cited, remixed, contested, taught, performed. The center can improve while the future dies.

In the paper’s terms: the center improves when SS is optimized; the future dies when g is not independently floored. The law says the floor must be exogenous; the coda says it must also be live; the conjunction — an exogenous injection channel that stays recombining — is the specification for the thing that needs building.

Appendix A — Derivation of the contraction condition

Let μ_t be a probability measure over a type space X , and let $\mu_{t+1} = (S \circ R) \mu_t$ for transmission kernel R and selection kernel S . Four axioms suffice for support contraction and entropy loss:

(A1) Finite sampling. R draws a finite sample from μ_t .

(A2) Non-expansion of support under R . $\operatorname{supp}(R \mu) \subseteq \operatorname{supp}(\mu)$. The transmission kernel does not introduce types absent from the current distribution.

(A3) Strict selection. S reweights toward a proper subset of $\operatorname{supp}(R \mu)$: there exists a non-negligible set A with $S(A) / (R \mu)(A) > 1$ and its complement down-weighted.

(A4) In-support stochasticity only. Any randomness in R (temperature, dropout, sampling noise) generates variation within $\operatorname{supp}(\mu)$, not outside it.

Under A1–A4: (i) $|\operatorname{supp}(\mu_{t+1})| \leq |\operatorname{supp}(\mu_t)|$ with probability 1, and strict inequality with positive probability whenever the support has more than one element. (ii) Shannon entropy $H(\mu_{t+1}) \leq H(\mu_t)$ when S is *monotonically type-concentrating* — reweighting that increases the mass on the already-heaviest types at the expense of lighter ones. Not all strict-selection operators satisfy this (a reweighting that equalizes mass across fewer types can raise entropy

over a reduced support); the condition on $\mathbb{S}$ is stronger than mere differential survival and should be verified for each substrate’s kernel. (iii) The mode concentrates: the probability mass on the highest-weight type increases monotonically.

These are standard properties of iterated resampling with selection; the contribution is their application to the semantic state, where A2 — no exogenous injection — is precisely the hinge the paper tests, and where violating A2 (injecting exogenous types at rate $\mathbb{g}_0 > 0$) is the case-1 floor that prevents exhaustion.

References

- Imported structure:- Allee, W. C. (1931). *Animal Aggregations*. Univ. Chicago Press.- Stephens, P. A., Sutherland, W. J., & Freckleton, R. P. (1999). *Oikos* 87(1): 185–190.- Courchamp, F., Clutton-Brock, T., & Grenfell, B. (1999). *TREE* 14(10): 405–410.- Courchamp, F., Berec, L., & Gascoigne, J. (2008). *Allee Effects in Ecology and Conservation*. Oxford Univ. Press.- Gilpin, M. E., & Soulé, M. E. (1986). In *Conservation Biology* (Soulé, ed.).- Lynch, M., Conery, J., & Burger, R. (1995). *Am. Nat.* 146: 489–518.- Eigen, M. (1971). *Naturwissenschaften* 58: 465–523.- Henrich, J. (2004). *American Antiquity* 69(2): 197–214.- Powell, A., Shennan, S., & Thomas, M. G. (2009). *Science* 324: 1298–1301.- Vaesen, K., et al. (2016). *PNAS* 113: E2241–E2247.- Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Univ. Chicago Press.- Mesoudi, A. (2011). *Cultural Evolution*. Univ. Chicago Press.
- Model collapse and linguistic flattening:- Shumailov, I., et al. (2024). *Nature* 631: 755–759.- Gerstgrasser, M., et al. (2024). *arXiv:2404.01413*.- Seddik, M. E. A., et al. (2024). *arXiv:2404.05090*.- Padmakumar, V., & He, H. (2024). *ICLR 2024*, *arXiv:2309.05196*.- Fitterer, S., Gangl, D., & Ulbrich, J. (2025). *ACL 2025 (SRW)*.- Doshi, A. R., & Hauser, O. P. (2024). *Science Advances* 10(28).
- Evidence bearing on the Ratchet:- Wan, A., & Kalman, M. (2025). *arXiv:2504.13868*.- Kleinberg, J., & Raghavan, M. (2021). *PNAS* 118(22): e2018340118.
- Evolutionary dynamics:- Page, K. M., & Nowak, M. A. (2002). *J. Theor. Biol.* 219(1): 93–98.- Sella, G., & Hirsh, A. E. (2005). *PNAS* 102(27): 9541–9546.- Rutherford, S. L., & Lindquist, S. (1998). *Nature* 396: 336–342.

Suggested Citation

Nobel Glas. “Diversity Contraction Across Substrates: A Boundary Law for Semantic Exhaustion” *Alexanarch*, 2026-06-02. <https://www.alexanarch.org/s/records/779/>

Deposit Information

This paper is deposit #779 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:02DD.EMPIRICAL` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/779/>
- License: CC BY 4.0
- Author ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)
- Provenance chain: alexanarch.org
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/779/>

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