Repair Pressure and the Lifecycle of Distinctions

Flyxion

Independent Researcher

Abstract

Distinctions are not merely represented; they are maintained. A distinction survives when some population continues paying the cost of preserving it, and disappears when no sustainable maintenance topology exists. We develop a dynamical theory of distinction maintenance in which collapsed distinctions evolve under three separable processes: failure recognition, repair generation, and repair coordination. The resulting phase space predicts five outcomes — stable collapse, migration, repair scarcity, repair competition, and canonical restoration — and explains why some distinctions disappear, some survive in specialist communities, and some return to dominant representations after extended periods of instability. Three results organize the argument. First, distinction migration is maintenance reallocation: a transition in maintenance topology rather than a change in semantic content. Second, repair pressure is driven by *attributed* failure rather than *objective* failure; the legibility operator separating these quantities explains why two systems can experience identical losses while exhibiting radically different repair dynamics. Third, canonical restoration requires repair generation and repair coordination as independent conditions, not merely sufficient repair pressure. Governance enters the framework not as an external consideration but as a determinant of the coordination coefficient governing variance reduction across competing repair candidates.

Contents

1. Introduction

Distinction migration is not primarily a semantic phenomenon but a maintenance-topology transition.

Consider technical vocabulary. When a distinction leaves ordinary speech and survives only among specialists — surgeons distinguishing types of tissue, pilots distinguishing approach clearances, ecologists distinguishing succession stages — the standard description says the distinction has been *preserved* in a specialist community. But this description conceals what actually happened. The distinction did not move because it retained meaning. It moved because a smaller population became willing to bear a larger per-capita maintenance cost. What changed was not the semantic content of the distinction but the distribution of who pays to keep it alive.

This observation, once made precise, has consequences well beyond lexical change. The same structure appears when a scientific community sustains a distinction that the broader culture has collapsed, when a bureaucratic institution preserves a category that no longer tracks the reality it was designed to partition, and when a machine learning system erases a distinction its training data did not preserve with sufficient frequency. In each case, the question that matters is not whether the distinction is meaningful but whether any population can sustain the maintenance cost of preserving it.

The present paper develops a dynamical theory built on this observation. The central objects are not distinctions themselves but the populations, institutions, and systems responsible for maintaining them. Lexical change, scientific revolutions, institutional drift, and machine learning failures appear as applications of a common framework rather than as separate phenomena requiring separate accounts.

Three contributions organize the argument. First, we formalize the claim that distinction migration is *maintenance reallocation*: a transition in maintenance topology rather than a change in semantic content. Second, we separate objective failure from attributed failure and show that repair pressure responds to the latter, not the former. This separation explains why two systems can suffer identical objective losses while exhibiting radically different repair dynamics. Third, we show that canonical restoration — the return of a collapsed distinction to the dominant representation — requires two independent conditions: a high-quality repair must become available, and coordination disagreement across the population must collapse. Repair pressure alone is insufficient for either.

The paper is structured as follows. Section 2 situates the framework in relation to four existing traditions. Section 3 develops the phase space and defines its regions. Section 4 derives the dynamical equations and identifies the three separable processes. Section 5 applies the framework to four cases spanning language, science, institutions, and machine learning. Section 6 examines what determines the coordination coefficient and connects the framework to the political economy of distinction governance. Section 7 concludes with a reinterpretation of paradigm shift and a set of open questions.

2. Background

Four traditions bear on the framework developed here, each contributing a partial account of the dynamics we formalize.

2.1. Lexical change and colexification

The comparative study of how languages partition semantic space has established that colexification — the merging of distinct meanings under a single lexical form — follows identifiable cross-linguistic patterns [3]. Work in this tradition has shown that conceptual proximity and frequency of co-occurrence predict which distinctions languages tend to collapse [4, 5]. What this tradition does not explain is why some distinctions resist collapse despite low frequency, why collapsed distinctions sometimes return, and why the same distinction can be maintained in one community while disappearing in another. The maintenance-topology framework addresses these gaps by replacing frequency-based accounts of survival with cost-bearing accounts.

2.2. Philosophy of science and paradigm change

Kuhn's account of scientific revolutions identified the phenomenon of theoretical turbulence: periods in which a community recognizes that the incumbent framework is failing without being able to agree on a replacement [1]. Lakatos extended this picture by showing that research programmes persist under empirical pressure through a protective belt of auxiliary hypotheses [2]. The framework developed here formalizes the turbulence regime as a region in a two-dimensional phase space and separates its two distinct mechanisms — repair scarcity, in which no adequate repair exists, and repair competition, in which adequate repairs exist but coordination fails. This separation generates different predictions about the dynamics of each regime. It also makes precise the observation, noted but not formalized by Kuhn, that paradigm adoption typically lags paradigm availability by a period determined by the coordination structure of the relevant community.

2.3. Institutional ecology and organizational theory

Research on organizational resilience has examined how institutions maintain distinctions — between categories, populations, procedures, and resource types — under environmental pressure [6, 7]. The concept of maintenance cost appears in this literature under various names: overhead, administrative burden, institutional friction [8]. What is less developed is a formal account of how maintenance costs are distributed across populations and how that distribution changes when institutions fail. The maintenance-topology concept formalizes this: what changes during institutional migration is not the distinction but the population bearing its cost.

2.4. Repair and resilience

Engineering resilience theory has long treated repair as a primary process rather than an exceptional one [11, 12]. The present framework adopts this orientation and applies it to representational systems: distinctions do not exist in stable states punctuated by occasional breakdown, but are continuously maintained against ongoing pressure toward collapse. This reframing has a methodological consequence. The appropriate question is not why some distinctions collapse but why any distinction survives, since all distinctions face continuous maintenance costs and only those with sustainable maintenance topologies persist.

3. The Phase Space

Let D be a distinction that has been collapsed, weakened, or removed from a representational system R held by a population Π operating in domain θ. The framework represents the lifecycle of D as motion through a two-dimensional phase space.

Remark 3.1 (Agent-relativity)**.** The phase space does not describe distinctions in isolation. It describes distinction-domain-population triples (D, θ, Π). The same distinction may occupy different regions of the space simultaneously for different populations, because repair pressure, maintenance cost, and repair quality are all population-relative quantities. The state of a distinction is indexed to the community responsible for maintaining it, not to the distinction itself. This is not a limitation of the framework but its central feature: it is what allows migration to appear as a structural phenomenon rather than an anomaly.

3.1. Axes

The **horizontal axis** is instantaneous repair pressure, defined as the net force on the system toward restoring distinction D at time t:

$$P_r(D, \theta, t) \;=\; R_D^\theta \cdot F(D, \theta, t) \cdot C(D, \theta, t) \;-\; M(D, \theta, t, \rho) \;-\; S(D, \theta, t), \qquad (1)$$

where R_D^θ is *reachability divergence* — the degree to which the collapsed distinction affects the set of futures accessible from the current state; F is the frequency with which the collapsed distinction is encountered; C is the consequence magnitude of failures attributable to the collapse; M is the maintenance cost of preserving the distinction under maintenance topology ρ; and S is *substitutability*, measuring the availability of workarounds that absorb failures without restoring the distinction.

Equation (1) defines P_r as a state variable: the instantaneous value at time t given the current configuration of costs, frequencies, and workarounds. Section 4 introduces the dynamics governing how P_r evolves over time; those dynamics are driven by attributed failure accumulation and substitution pressure, not by Equation (1) directly. The two formulations are complementary: the state equation gives the value; the dynamical equation gives the rate of change.

The **vertical axis** is available repair quality:

$$Q^*(D, \theta, t) \;=\; \max_{R \in \mathcal{R}_t} \; \mathbb{E}_{\pi \sim \Pi}[Q(R \mid \pi)]\,, \qquad (2)$$

where $\mathcal{R}_t$ is the set of repairs available at time t and

$$Q(R \mid \pi) \;=\; \frac{\Delta\,\mathrm{Vol}(\mathcal{A})}{\Delta K(R \mid \pi)} \tag{3}$$

is the efficiency of repair R relative to prior π. Here $\Delta\,\mathrm{Vol}(\mathcal{A})$ is the recovered admissible future volume and $\Delta K(R \mid \pi)$ is *adoption complexity*: the cost of integrating R into prior π, measured not as absolute representational complexity but as the cost of restructuring existing representational commitments. A repair that is compact in absolute terms can carry high $\Delta K(R \mid \pi)$ for a population with deep investment in the incumbent representation.

Proposition 3.2 (Canonical Repair Criterion)**.** A repair R need not maximize $Q(R \mid \pi)$ for any individual prior π in order to become canonical. It is sufficient that it maximize $\mathbb{E}_{\pi\sim\Pi}[Q(R \mid \pi)]$ across the population.

Proof sketch. Region IV is defined by the condition $\mathbb{E}_{\pi\sim\Pi}[Q(R^* \mid \pi)] \geq \tau_Q$, which is a condition on expected cross-prior quality rather than on prior-specific quality. A repair that achieves moderate efficiency across many priors may therefore satisfy this condition while a repair that achieves very high efficiency for a narrow subset of priors does not. The canonical repair is the one that clears the cross-prior adoption threshold, not the one that is optimal for any particular prior. □

This result explains why repairs that are highly persuasive within one subfield often fail to become canonical while less technically impressive repairs that address a wider range of prior commitments succeed. Plate tectonics became canonical not because it was the most compelling account for any single geological prior, but because seafloor spreading raised $Q(R \mid \pi)$ across geophysical, paleontological, and biogeographical priors simultaneously.

Remark 3.3 (ΔK hysteresis)**.** Adoption complexity exhibits strong path-dependence. Maintaining a distinction under duress — when the incumbent representation is still intact but under pressure — is substantially cheaper than restoring a distinction after its structural hooks have been removed from the environment. Once a distinction has been purged from codebases, regulatory forms, training curricula, and professional standards, $\Delta K(R \mid \pi)$ for any restoring repair reflects not just cognitive restructuring but capital replacement: legacy systems must be rewritten, printed materials revised, and credentialing pipelines rebuilt. This asymmetry means that the phase-space boundary between Region I and the restoration path is not symmetric with the boundary at which collapse originally occurred. A distinction may collapse at a relatively low maintenance cost threshold and require a much higher repair pressure to restore, because ΔK has grown substantially during the interval of collapse. Hysteresis in ΔK is therefore a mechanism by which Region I becomes increasingly stable over time even when the underlying reachability divergence R_D^θ is unchanged.

Proposition 3.4 (Restoration Hysteresis)**.** Suppose a distinction D collapses at time t_0 and remains collapsed over the interval $[t_0, t_1]$. If adoption complexity satisfies

$$\frac{d}{dt}\Delta K(R \mid \pi) \;\geq\; 0$$

during the collapse interval, then the repair pressure required for restoration at t_1 exceeds the repair pressure that would have been required to preserve the distinction at t_0:

$$P_r^{\text{restore}}(t_1) \;>\; P_r^{\text{preserve}}(t_0).$$

Proof sketch. Restoration requires integrating the distinction into a representational system from which supporting structures have already been removed. Since ΔK appears in the denominator of repair quality,

$$Q(R \mid \pi) \;=\; \frac{\Delta\,\text{Vol}(\mathcal{A})}{\Delta K(R \mid \pi)},$$

an increase in ΔK reduces $Q(R \mid \pi)$ for fixed recovered admissible volume. The adoption threshold τ_Q is therefore harder to reach at t_1 than at t_0, requiring higher repair pressure to achieve the same restoration outcome. The collapse and restoration boundaries in the phase space are consequently non-symmetric. □

The proposition gives a mathematically precise grounding for the informal observation that collapsed distinctions become harder to restore the longer they remain collapsed. The asymmetry is not a contingent feature of any particular case but a structural consequence of ΔK's role in the denominator of repair quality.

A third quantity, **coordination disagreement** $\text{Var}_\pi(Q)$, is not an axis but governs which sub-region of the high-P_r half of the space a system occupies.

3.2. Agent-relativity: a worked example

The agent-relativity remark in Section 3 is not a qualification appended to an otherwise universal theory. It is a constitutive feature of the framework. The phase space does not describe where a distinction is. It describes where a distinction is *for a given population*. This subsection works through a concrete case to show why the distinction matters and what would be lost by ignoring it.

Consider the automation-augmentation distinction as it stands in contemporary AI development — the distinction between deploying AI to substitute for human labor and deploying it to extend the range of what humans and machines can accomplish together. Ask where this distinction sits in the phase space.

The question has no answer without specifying a population.

For the population of corporate accounting regimes that measure AI return on investment primarily through headcount reduction, the distinction occupies **Region I**. Repair pressure is low not because the distinction lacks reachability divergence — it has enormous reachability divergence, since the two pathways lead to substantially different future volumes — but because the legibility operator $\mathcal{L}_\delta$ embedded in the measurement regime approaches zero for augmentation value. Failures attributable to the collapsed distinction accumulate as noise. No attributed pressure is generated. The distinction is stable in collapse for this population, and the stability is self-reinforcing through the parasitic substitution mechanism: workarounds absorb failure signatures and prevent attribution from being established.

For the population of researchers in human-computer interaction, cognitive science, and labor economics who study complementarity between human and machine capability, the same distinction simultaneously occupies **Region III-b**. Repair pressure is high — the community is acutely aware that the distinction is being collapsed in industry and regards this as a significant failure. A variety of high-quality repairs are available: augmentation-centered design frameworks, capability-preserving deployment protocols, alternative accounting methodologies. But coordination disagreement is also high: different subpopulations within this research community rank the available repairs differently depending on their theoretical priors, and no single repair has achieved sufficient cross-prior quality to become canonical within the community, let alone to propagate outward to the accounting regimes in Region I.

The same distinction, at the same time: Region I for one population, Region III-b for another. This is not a paradox. It is what the agent-relative framework predicts, and it is what a population-independent reading of the phase space would obscure.

The case also illustrates a dynamical implication. The Region III-b population cannot directly repair the distinction for the Region I population by raising their own Q^* or reducing their own $\mathrm{Var}_\pi(Q)$. What would need to change is either the legibility operator operating within the accounting regime — a change to the measurement infrastructure itself — or the repair pressure experienced by the Region I population, which would require either a rise in attributed failure or the degradation of the substitution mechanisms currently keeping attributed pressure near zero. These are different interventions, targeting different populations, and the framework makes the difference precise.

The critical repair pressure threshold τ_P is not a free parameter. We set

$$\tau_P \;=\; M_{\text{general}}(D, \theta, t, \rho), \tag{4}$$

making the primary boundary structurally determined by maintenance cost rather than stipulated externally. The migration inequality therefore becomes

$$M_{\text{specialist}} \;<\; P_r \;<\; M_{\text{general}}, \tag{5}$$

and the Region I/II boundary coincides with the lower bound of this interval.

The critical repair quality threshold τ_Q is the minimum expected cross-prior quality required for a repair to overcome adoption inertia in the relevant population.

Proposition 3.5 (Migration Criterion)**.** A distinction D in domain θ enters Region II if and only if repair pressure exceeds specialist maintenance cost while remaining below general maintenance cost:

$$M_{\text{specialist}}(D, \theta, t, \rho) \;<\; P_r(D, \theta, t) \;<\; M_{\text{general}}(D, \theta, t, \rho).$$

Proof sketch. Region I requires $P_r \;<\; \tau_P \;=\; M_{\text{general}}$, so the distinction carries insufficient pressure for general-population maintenance. Region III requires $P_r \geq \tau_P = M_{\text{general}}$, so the distinction carries pressure sufficient for general maintenance but faces a repair-quality or coordination obstacle. Region II is precisely the intermediate band: pressure is sufficient to justify maintenance for a subpopulation with lower per-capita cost $M_{\text{specialist}} < M_{\text{general}}$, but insufficient to sustain general-population maintenance. A specialist subpopulation exists if and only if such a cost differential obtains; their willingness to bear higher per-capita cost is what makes the migration stable. The bounds of the inequality are therefore both necessary and sufficient for Region II occupation. □

The proposition makes precise the intuition that migration is not a degenerate form of collapse or restoration but an equilibrium state supported by a specific cost structure. It also entails that migration is reversible in both directions: if P_r falls below $M_{\text{specialist}}$, the specialist community can no longer sustain maintenance and the distinction dies; if P_r rises above M_{general}, general-population maintenance becomes viable and the distinction may return to broader use.

Proposition 3.6 (Migration Stability)**.** Assume $M_{\text{specialist}} < P_r < M_{\text{general}}$. Then Region II is dynamically stable under perturbations that do not alter the maintenance-cost ordering.

Proof sketch. Small increases in P_r remain insufficient to justify general-population maintenance as long as $P_r < M_{\text{general}}$. Small decreases in P_r remain sufficient to justify specialist maintenance as long as $P_r > M_{\text{specialist}}$. Consequently the distinction remains preserved within the specialist topology under small perturbations. Exit from Region II requires crossing one of the maintenance boundaries rather than merely perturbing the state within them. □

The stability result means that migration is not a transitional state on the way to either death or restoration. It is a genuine equilibrium regime that can persist indefinitely as long as the maintenance-cost ordering is preserved. Many specialist distinctions that appear to be "holding on" are in fact in a dynamically stable configuration rather than one that is slowly collapsing.

3.3. Regions

Table 1. Summary of phase-space regions.

Region	Label	Conditions
I	Stable Collapse	$P_r < \tau_P$
II	Migration	$M_{\text{spec}} < P_r < M_{\text{gen}}$
III-a	Repair Scarcity	$P_r \geq \tau_P;\ \max_R \mathbb{E}_\pi[Q(R\|\pi)] < \tau_Q$
III-b	Repair Competition	$P_r \geq \tau_P;\ \max_R \mathbb{E}_\pi[Q] \geq \tau_Q;$ $\text{Var}_\pi(Q) \gg 0$
IV	Canonical Restoration	$P_r \geq \tau_P;\ \mathbb{E}_\pi[Q(R^*\|\pi)] \geq \tau_Q;\ \text{Var}_\pi(Q)$ below coordination threshold

Region I — Stable Collapse. $P_r < \tau_P$. Repair pressure does not exceed general maintenance cost. The distinction remains collapsed. Failures are rare, low-consequence, absorbed by substitution, or — crucially — illegible as attributable to the collapse. A distinction can impose continuous objective costs and still remain in Region I indefinitely if the connection between those costs and the collapsed distinction is never established. This is the regime of silent institutional failure and unrecognized classification error.

Region II — Migration. $M_{\text{specialist}} < P_r < M_{\text{general}}$. Repair pressure exceeds the maintenance cost for a specialist subpopulation but not for the general population. The distinction survives by changing its maintenance topology: the burden of preservation is transferred from the general system to a smaller, more invested community willing to bear higher per-capita cost. Migration is maintenance reallocation. The distinction that survives in technical vocabulary, expert jargon, or professional classification has not merely been semantically preserved — it has found a population for whom the maintenance cost is worth paying.

Region III-a — Repair Scarcity. $P_r \geq \tau_P$ and $\max_R \mathbb{E}_\pi[Q(R \mid \pi)] < \tau_Q$. The community recognizes that collapse is costly but no available repair achieves sufficient cross-prior quality. Candidate repairs proliferate and fail. This is the regime of failed precursors: the system is searching for repairs because none of the available candidates work well enough. The constraint is not motivation but adjacent-possible availability.

Region III-b — Repair Competition. $P_r \geq \tau_P$ and $\max_R \mathbb{E}_\pi[Q(R \mid \pi)] \geq \tau_Q$ but $\mathrm{Var}_\pi(Q)$ exceeds the coordination threshold. A high-quality repair exists for some priors. The obstacle is not invention but synchronization. Different populations rank candidate repairs differently. The limiting dynamics are coalition formation, persuasion, institutional authority, and demographic replacement rather than theoretical innovation. This regime is distinguishable from III-a by the pattern of activity it generates: in III-a communities produce experimental frameworks; in III-b they produce competing schools.

Region IV — Canonical Restoration. $P_r \geq \tau_P$, $\mathbb{E}_\pi[Q(R^* \mid \pi)] \geq \tau_Q$, and $\mathrm{Var}_\pi(Q)$ has collapsed below the coordination threshold. A repair achieves sufficient cross-prior quality to overcome adoption inertia across the relevant population. The distinction is restored and stabilizes in the dominant representation. Canonical repairs are not simply correct repairs. They are repairs whose efficiency exceeds the adoption threshold across a sufficiently broad distribution of priors held by the relevant population.

3.4. Phase diagram

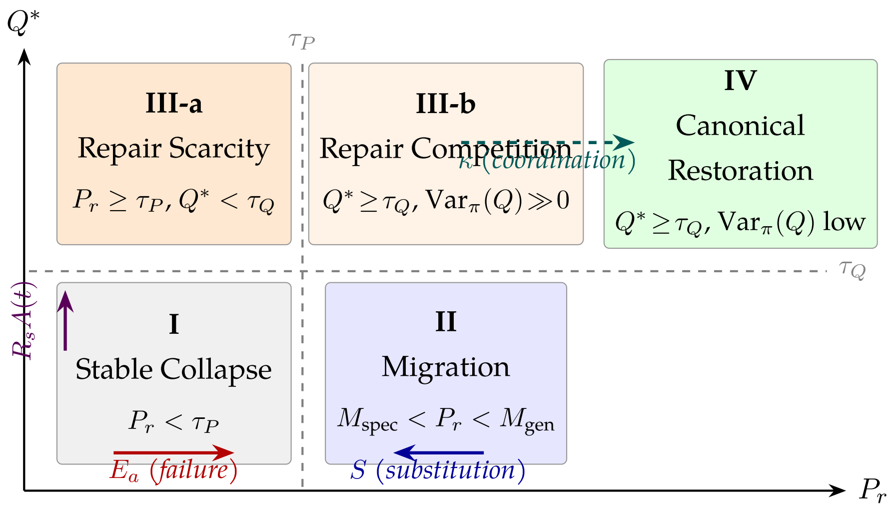

Figure 1. Phase space of distinction maintenance dynamics. Horizontal axis: repair pressure P_r (Eq. 1). Vertical axis: available repair quality Q^* (Eq. 2). Regions: **I** stable collapse ($P_r < \tau_P$); **II** migration ($M_{\text{spec}} < P_r < M_{\text{gen}}$); **III-a** repair scarcity (high P_r, low Q^*); **III-b** repair competition (high P_r, high Q^*, high $\mathrm{Var}_\pi(Q)$); **IV** canonical restoration (high P_r, high Q^*, low $\mathrm{Var}_\pi(Q)$). *Directional forces:* attributed failure E_a pushes rightward ($\rightarrow$); substitution S pushes leftward ($\leftarrow$); repair effort combined with adjacent-possible availability $R_s \cdot A(t)$ pushes upward ($\uparrow$); coordination coefficient κ reduces $\mathrm{Var}_\pi(Q)$ and drives the III-b $\rightarrow$ IV transition. The diagram describes distinction-domain-population triples (D, θ, Π), not distinctions in isolation; the same distinction may occupy different regions for different populations simultaneously.

4. Dynamics

The phase space in Section 3 is static. To explain how distinctions move between regions we require equations of motion. Three separable processes govern the dynamics: failure recognition, repair generation, and repair coordination.

4.1. Failure recognition

$$\frac{dP_r}{dt} = \alpha\, E_a(t) - \beta\, S(t), \tag{6}$$

where $E_a(t) = \mathcal{L}_\delta(E_f)$ is the *attributed failure process* and $\mathcal{L}_\delta$ is a legibility operator incorporating attribution delay δ.

Equation (6) governs the evolution of the state variable defined in Equation (1). The two formulations are consistent: the state equation gives P_r as a function of current costs and frequencies, while the dynamical equation gives its rate of change as a function of failure accumulation and substitution. Together they determine both the position of a distinction in the phase space and its trajectory.

The separation of E_f (objective failures) from E_a (attributed failures) is the first non-obvious move in the dynamics. Repair pressure responds to attributed failure, not to objective failure. Two systems can suffer identical objective losses and exhibit radically different repair dynamics because one has established the causal connection between the failure and the collapsed distinction while the other has not.

The attribution delay δ is not a fixed constant but a function of the failure mode's visibility structure. Failures that produce discrete, attributable events — crashes, deaths, measurable errors — have short δ. Failures that produce diffuse, chronic, multi-causal degradation have long δ and may never generate sufficient attributed pressure to enter Region III. This explains a class of collapsed distinctions that impose continuous objective costs without ever generating repair dynamics: their failure mode is structurally illegible, and $\mathcal{L}_\delta$ approaches zero regardless of the magnitude of E_f.

Substitution mechanisms S provide counterpressure, pushing systems leftward in the phase space. Workarounds, informal circumlocutions, bypass procedures, and specialist knowledge that compensates for the collapsed distinction all reduce effective P_r. Substitution mechanisms are themselves subject to maintenance costs, however, and typically degrade over time as the population bearing their cost shrinks or as the workarounds become more elaborate than the original distinction would have been.

Remark 4.1 (Parasitic substitution trap)**.** A latent feedback between S and $\mathcal{L}_\delta$ can produce a Region I attractor from which exit becomes structurally blocked. When substitution mechanisms become sufficiently institutionalized, they do not merely reduce P_r — they actively conceal the original collapse by absorbing its failure signature into the workaround infrastructure. Failures that would otherwise accumulate as attributed pressure on the collapsed distinction are instead attributed to the increased complexity or occasional breakdown of the substitution mechanism itself. The result is that high S can drive $\mathcal{L}_\delta \to 0$: objective failures E_f continue to occur and may even scale as the workarounds proliferate, but the causal connection to the collapsed distinction is never established, and E_a remains near zero. A system in this configuration can sustain arbitrarily large objective failure accumulation while generating no repair pressure, because the attribution structure has been permanently redirected. The appropriate diagnostic for this condition is not measurement of P_r — which will appear low — but independent measurement of E_f against the predicted failure rate under an intact distinction, compared to the observed rate under the substitution regime.

Proposition 4.2 (Legibility Bound)**.** If $\mathcal{L}_\delta(E_f) \to 0$, then

$$\frac{dP_r}{dt} \to -\beta S.$$

Consequently a distinction may generate unbounded objective failure while repair pressure remains bounded above by zero.

Proof. Substituting $E_a = \mathcal{L}_\delta(E_f)$ into $\frac{dP_r}{dt} = \alpha E_a - \beta S$ and taking the limit as $\mathcal{L}_\delta(E_f) \to 0$ yields $\frac{dP_r}{dt} \to -\beta S$ directly. Since $S \geq 0$ and $\beta > 0$, the right-hand side is non-positive. Repair pressure therefore cannot increase and may decrease without bound on the substitution term, regardless of the magnitude of E_f. □

The proposition gives a formal statement of the permanent Region I phenomenon. It establishes that illegibility is not merely a delay in the generation of repair pressure but a structural ceiling: when the legibility operator is zero, repair pressure is governed entirely by substitution dynamics and is insensitive to objective failure accumulation.

4.2. Repair generation

$$\frac{dQ^*}{dt} = \gamma R_s \cdot A(t), \tag{7}$$

where R_s is repair-seeking effort and $A(t)$ is *adjacent-possible availability*.

The separation of effort from opportunity explains a pattern that appears repeatedly in the history of science and institutions: long periods of intense repair-seeking with little progress, followed by rapid improvement. If repair generation depended only on effort, we would predict steady improvement whenever motivated communities are working on a problem. The historical record shows otherwise. Communities can be highly motivated and still trapped in Region III-a for extended periods because $A(t)$ is low — the conceptual, technological, or institutional materials required for a high-quality repair do not yet exist.

Adjacent-possible availability $A(t)$ is not equivalent to technological availability, though technology is one of its determinants. The concept draws on the notion of the adjacent possible as the space of first-order transformations available from the current state [9, 10]. $A(t)$ increases when new measurement techniques make previously invisible distinctions tractable, when conceptual innovations imported from adjacent domains make new repair candidates thinkable, when institutional reforms make previously uncoordinated repair efforts possible, and when accumulated failure experience makes the failure pattern legible in ways that suggest repair directions. Microscopy raised $A(t)$ for germ theory. Ocean-floor mapping raised $A(t)$ for plate tectonics.

4.3. Repair coordination

$$\frac{d}{dt}\,\mathrm{Var}_{\pi}(Q) \;=\; -\kappa, \tag{8}$$

where κ is the *coordination coefficient* governing the rate at which disagreement about repair quality collapses across the population.

Equation (8) captures the process that distinguishes Region III-b from Region IV. A high-quality repair can exist — $Q^* \geq \tau_Q$ — while coordination disagreement remains high enough to prevent adoption. The transition to Region IV requires $\mathrm{Var}_{\pi}(Q)$ to fall below a threshold, and the rate of that fall is determined by κ.

κ is high when institutional authority can impose a repair on agents with high adoption complexity $\Delta K(R \mid \pi)$, when prestige structures are concentrated enough that a small number of prior-holders can shift the distribution, or when generational turnover is rapid relative to the turbulence timescale. Generational turnover is a particularly reliable variance-reduction mechanism: it removes the agents with the highest $\Delta K(R \mid \pi)$, since they carry the largest prior investment in the incumbent representation. The lag between paradigm availability and paradigm adoption that appears throughout the history of science is largely a κ-determined phenomenon.

κ is low when priors are entrenched in material infrastructure that is expensive to replace regardless of persuasion, when the relevant population is fragmented across non-communicating communities, or when variance in $Q(R \mid \pi)$ across priors is driven by political interest rather than epistemic disagreement. In the latter case, agents are not failing to converge on the best repair — they are defending positions. This condition can maintain high $\mathrm{Var}_\pi(Q)$ indefinitely even when a high-quality repair is broadly recognized, because adoption would require conceding interests rather than updating beliefs. Governance enters the framework at exactly this point, as discussed in Section 6.

Proposition 4.3 (Coordination Bottleneck)**.** A necessary condition for transition from Region III-b to Region IV is

$$\lim_{t \to t^*} \mathrm{Var}_\pi(Q) \ < \ \tau_{\text{coord}},$$

regardless of the magnitude of Q^*.

Proof. Region IV requires both $Q^* \geq \tau_Q$ and $\mathrm{Var}_\pi(Q) < \tau_{\text{coord}}$ by definition. The first condition does not imply the second: Q^* may be arbitrarily large while $\mathrm{Var}_\pi(Q)$ remains above the coordination threshold. Therefore high repair quality is not sufficient for canonical restoration, and the variance condition is a genuinely independent necessary condition. □

The proposition formalizes one of the framework's central claims: invention and adoption are separable processes with independent necessary conditions. A community may possess a repair of arbitrarily high quality and remain in turbulence indefinitely if coordination disagreement is not resolved. This is not a contingent sociological observation but a structural feature of the phase space.

4.4. Default trajectory

Remark 4.4 (Population-indexed dynamics)**.** The dynamical system is population-indexed. More precisely, dP_r/dt, dQ^*/dt, and $d\ \mathrm{Var}_\pi(Q)/dt$ should be understood as $dP_r^{(\Pi)}/dt$, $dQ^{*(\Pi)}/dt$, and $d\ \mathrm{Var}_\pi^{(\Pi)}(Q)/dt$, with the population index Π suppressed throughout for readability. Distinct populations may therefore instantiate distinct trajectories for the same distinction-domain pair (D, θ). The framework does not predict a unique lifecycle for a distinction. It predicts a *family* of population-relative lifecycles whose interactions generate the migration, restoration, and persistence phenomena described in the preceding sections.

This indexing separates the present framework from accounts in which a distinction or paradigm has a single state at a given time — in Kuhnian terminology, a shared state across the relevant scientific community. The community is not an atomic unit here. It is a distribution over priors $\pi \sim \Pi$, and different subpopulations within it may be in different regions of the phase space simultaneously. What looks from the outside like a community "in turbulence" may be a superposition of a Region I subpopulation (where substitution mechanisms have suppressed attributed pressure), a Region III-a subpopulation (actively searching for repairs), and a Region III-b subpopulation (holding competing high-quality repairs that cannot coordinate). These conditions require different interventions, and treating the community as a unit causes the interventions to interfere.

In the absence of active repair-seeking, P_r is non-decreasing for any collapsed distinction as failure events accumulate subject to substitution pressure and attribution delay. The default trajectory moves rightward in the phase space, crossing τ_P when accumulated attributed pressure exceeds general maintenance cost. Systems then enter Region III-a if $A(t)$ is low, or Region III-b if a high-quality repair exists but $\mathrm{Var}_\pi(Q)$ remains elevated. Exit into Region IV requires both $Q^* \geq \tau_Q$ and $\mathrm{Var}_\pi(Q)$ below the coordination threshold — conditions that are independent and may be satisfied at different times, producing the characteristic lag between repair availability and repair adoption.

5. Applications

We examine four cases chosen to span the region structure of the phase space and to demonstrate that the framework applies across domains without requiring domain-specific modification.

5.1. Legal distinction migration: Region I to Region II

Legal systems offer among the clearest examples of Region II dynamics because the maintenance cost structure is explicit and institutionally enforced. The common-law distinction between *trespass vi et armis* and *trespass on the case* — the former requiring direct forcible injury, the latter covering indirect harm — was central to English tort law for centuries. As the distinction migrated out of ordinary legal practice through procedural reforms in the nineteenth century, it did not disappear. It survived in legal scholarship, in specialist doctrinal analysis, and in the case law of jurisdictions that retained older procedural frameworks.

The maintenance-topology shift is precise. The distinction moved from a distinction that ordinary practitioners had to track — enforced by pleading rules that punished misdescription — to a distinction maintained by legal historians, comparativists, and scholars of juridical theory. The per-capita cost of maintenance rose; the population bearing that cost shrank. The distinction was not semantically lost. Its maintenance topology changed.

The structural point generalizes. Many legal distinctions survive only in specialist subpopulations long after their operational role in ordinary practice has been eliminated. In each case, the persistence is not explained by the distinction's continued semantic value to the general legal community but by the existence of a subpopulation — academic, jurisdictional, or historical — for whom the maintenance cost remains worth paying.

5.2. Semmelweis and puerperal fever: Region III-a

Ignaz Semmelweis demonstrated in the 1840s that handwashing dramatically reduced puerperal fever mortality in obstetric wards. His intervention worked. It did not become canonical.

The framework identifies this as a Region III-a outcome. Repair pressure was high: the mortality rates in physician-attended wards were well-documented and recognized as a failure. Semmelweis offered a repair that demonstrably reduced failures in the immediate context. But his repair was a local intervention — it reduced mortality without providing an explanation that integrated the failure pattern into the existing representational structure of medicine. $\Delta K(R_{\text{Semmelweis}} \mid \pi_{\text{medical}})$ was high for the medical community because accepting his intervention without explanation required abandoning or suspending large portions of the incumbent theoretical framework. The repair had high local efficiency but low cross-prior Q.

The attribution-lag mechanism is also visible here. Many of the deaths attributable to the collapsed distinction between contaminated and uncontaminated contact were not attributed to that collapse. They were attributed to miasma, constitutional predisposition, and other categories that the incumbent framework made available. $\mathcal{L}_\delta$ was low: the failure mode was visible but the attribution to the relevant distinction was not established in the wider community. Semmelweis had established it locally; the general system had not.

Germ theory provided the globally admissible repair. It did not merely explain puerperal fever. It provided a mechanism that made the entire family of failures — wound infection, surgical mortality, epidemic disease, food spoilage

— simultaneously legible as manifestations of the same collapsed distinction. $\Delta\,\mathrm{Vol}(\mathcal{A})$ was large because the repair recovered admissible futures across many previously disconnected failure domains. $\Delta K(R_{\text{germ}} \mid \pi)$ remained high initially but fell as $A(t)$ rose with microscopy and as generational turnover reduced the population bearing the highest prior investment in miasma theory.

5.3. Plate tectonics: Region III-b to Region IV

Alfred Wegener proposed continental drift in 1912. It was not widely accepted until the 1960s. The intervening five decades are standardly described as a period of resistance or conservatism. The framework offers a more precise description.

Wegener's proposal entered a field in which the incumbent framework — contractionism, the hypothesis that geological features were produced by the cooling and contraction of a once-molten Earth — was failing to account for an accumulating set of anomalies: the fit of continental coastlines, the distribution of fossil species across ocean basins, the similarity of geological formations on separated continents. Repair pressure was high. The contractionist framework was generating recognized failures.

The obstacle was not that Wegener's repair was unknown. It was that $Q(R_{\text{Wegener}} \mid \pi)$ varied enormously across the geological community depending on prior commitments. For geologists whose priors included close attention to paleontological and biogeographical evidence, Q was high. For geologists whose priors centered on geophysical mechanism, Q was low because Wegener could not specify a plausible driving force for continental movement. $\mathrm{Var}_{\pi}(Q)$ was high. The system was in Region III-b.

The transition to Region IV occurred when ocean-floor mapping in the 1950s raised $A(t)$ by making mid-ocean ridges and their symmetrical magnetic striping accessible as evidence. This raised $Q(R_{\text{plate tectonics}} \mid \pi)$ for the geophysical prior by providing a mechanism — seafloor spreading — that addressed the driving-force objection. The cross-prior variance collapsed because the new evidence was interpretable within multiple prior frameworks and raised Q for each of them. κ rose as a younger generation without prior investment in contractionism entered the field. The transition to Region IV was rapid once both conditions were met.

The case also clarifies the distinction between Region III-a and Region III-b.

The Semmelweis case was III-a: the community needed a repair that did not yet exist. The Wegener case was III-b: the repair existed but could not coordinate adoption. These are different diagnoses with different implications. In III-a, the intervention that matters is raising $A(t)$. In III-b, the intervention that matters is raising κ.

5.4. Use-mention collapse in large language models: permanent Region I

In logic and linguistics, the distinction between *using* a word to refer to something and *mentioning* a word as a linguistic object is marked by convention and consequential in reasoning. The sentence *snow is white* uses the word *snow*; the sentence *'snow' has four letters* mentions it. Collapsing this distinction produces systematic errors in reasoning about language, meaning, and reference.

Large language models trained on natural language corpora systematically collapse this distinction. Training data does not preserve the use-mention distinction with sufficient frequency or consequence to generate repair pressure within the training objective. The collapsed distinction produces failures — errors in reasoning about quotation, reference, metalinguistic statements, and certain classes of logical inference — but the failure mode carries high δ. The connection between the errors and the collapsed distinction is not established within the system's error signal. $\mathcal{L}_\delta$ approaches zero: the failures accumulate as noise rather than as attributed pressure on a specific collapsed distinction.

This is a case of permanent Region I via illegibility. The distinction is not in Region II because no specialist subpopulation within the training system bears the maintenance cost. It is not in Region III because attributed repair pressure does not rise. The system continues to fail in the relevant ways while the failure mode remains unrecognized as a maintenance-topology problem.

The case demonstrates that illegibility is not merely a delay. It is a structural condition that can hold a system in Region I indefinitely regardless of objective failure accumulation. The appropriate intervention is not increased training on the same data but a change in the attribution structure: a modification to what counts as a legible failure within the training objective.

5.5. The automation–augmentation distinction: Region I via measurement compression

A second class of distinction collapse in AI systems operates at the level of design objectives rather than training data. The distinction between *automation* — deploying AI to substitute for human labor — and *augmentation* — deploying AI to extend the range of what humans and machines can accomplish together — has high reachability divergence. Which design philosophy a firm or institution pursues determines substantially different futures: different distributions of productive capacity, different distributions of who captures economic gains, and different trajectories for what kinds of problems become tractable at all. A world in which AI is primarily used to do what humans already do at lower cost reaches different futures than one in which AI is used to do what neither humans nor machines could previously do.

The distinction is collapsing not because it is semantically unclear but because the dominant measurement regime makes it illegible as a governance failure. When the return on AI investment is measured primarily as headcount reduction — labor costs avoided — the augmentation pathway generates no positive signal within the accounting structure even when it creates substantially more value. The legibility operator $\mathcal{L}_\delta$ for "we are foregoing augmentation value" approaches zero: the failure mode does not register as a failure within the metrics that drive institutional decision-making, because the value of unrealized augmentation never appears in the ledger. Objective losses accumulate — in foregone discovery, in products and services that do not exist, in human capabilities that atrophy rather than develop — but attributed pressure on the collapsed distinction remains near zero.

This is a Region I case with the same structure as the use-mention collapse in language models: a measurement regime functions as a legibility operator that systematically misattributes failures. The difference is that here the measurement regime is not a training objective but a corporate accounting convention, and the population that would bear the maintenance cost of the augmentation distinction — workers, researchers, and institutions invested in the complementarity of human and machine capability — is precisely the population whose bargaining position is weakened by the automation pathway the measurement regime favors.

The reachability divergence R_D^θ between the two pathways is therefore not visible to the agents controlling the compression scheme, because the futures foreclosed by automation are futures in which those agents would have had less concentrated control over productive capacity. The measurement compression and the governance compression reinforce each other: the distinction stays collapsed because the attribution structure that would make its collapse legible is itself controlled by the agents who benefit from the collapse.

6. Distinction Governance

The coordination coefficient κ is not exogenous to the systems the framework describes. It is produced by institutions, professions, educational systems, regulatory structures, and authority relationships that function partly as κ-raising mechanisms. Understanding what determines κ is therefore equivalent to understanding the political economy of distinction maintenance.

6.1. Power as compression

Every representational system compresses reality into categories, and every compression scheme preserves some distinctions while collapsing others. The choice of what gets merged and what remains distinct determines which futures remain navigable for which agents. In this sense, the ability to impose a compression scheme on other agents is a form of power: it sets the maintenance topology that others must work within or work around. Bureaucracies do not merely record categories. They determine which categories are operationally real, which failures are legible, and which repairs are institutionally possible.

Compression is necessary rather than optional: no agent can maintain every distinction, and all representational systems must allocate maintenance costs across a finite budget. The governance question is therefore not whether compression occurs but how the resulting maintenance costs are distributed — which distinctions bear the cost of being collapsed, and which populations absorb the cost of workarounds when they do.

6.2. Governance and κ

The institutions that raise κ are also the institutions that determine which repairs get coordinated. A professional accreditation body that standardizes terminology raises κ for distinctions within its domain while potentially lowering $A(t)$ for repairs that would require terminology revision. A regulatory agency that mandates categories raises κ for the mandated distinctions while imposing high $\Delta K(R \mid \pi)$ on any repair that would require reconceptualizing the regulatory framework. A university curriculum that transmits standard distinctions raises κ for students who share the curriculum while generating prior-variance between graduates and practitioners who acquired distinctions through different channels.

In each case, the governance structure shapes not only which repairs get adopted but which repairs become representable as candidates. This is the mechanism by which governance enters the phase dynamics: not by changing P_r directly, but by controlling the rate at which $\text{Var}_\pi(Q)$ contracts and the range of repairs for which $\Delta K(R \mid \pi)$ is tractable.

6.3. Migration as maintenance reallocation under imposed compression

When a governance structure imposes a compression that collapses distinctions a subpopulation continues to need, the maintenance burden reallocates. That subpopulation develops workarounds, informal classifications, specialist vocabularies, and practical distinctions that the official representation does not carry. From the perspective of the phase space, this is a migration event: the distinction moves from a maintenance topology supported by the general system to one supported by a smaller community bearing higher per-capita cost.

The formal consequence is that imposed compression does not eliminate maintenance costs — it redistributes them. Populations whose distinctions are collapsed by an official scheme absorb the ongoing cost of preservation through informal channels. When those informal channels degrade — through community dispersal, generational discontinuity, or resource constraint — the maintenance topology fails and the distinction dies. When they remain robust, the distinction persists in Region II indefinitely, invisible to the official scheme but operationally present in practice.

The framework therefore identifies a structural consequence of imposed compression: the gap between the official representation and the operational distinctions actually in use. That gap is not a failure of implementation; it is the predictable outcome of a maintenance-cost redistribution that the official scheme does not account for.

6.4. Compression at scale and the conditions for distributed maintenance

The classical argument for distributed economic coordination rests on the observation that knowledge is local, dispersed, and impossible for any central agent to aggregate completely. On this account, the superiority of market mechanisms over central planning follows from an epistemic constraint: no single compression scheme can capture enough of the distinctions that matter to coordinate production efficiently across an entire economy.

The emergence of systems capable of aggregating and processing information at previously unavailable scale changes the conditions under which this argument holds. When a sufficiently powerful compression scheme can be imposed across an entire domain — standardizing the categories through which transactions, labor contributions, and productive activities are recorded and valued — the maintenance cost of distinctions that fall outside the scheme is no longer distributed across many local agents. It is concentrated in the populations whose activities the scheme fails to capture. The question of who bears that cost, and whether any maintenance topology can sustain the distinctions the scheme collapses, becomes a governance question of the first order.

The framework makes this precise: when a single compression scheme achieves sufficient κ-raising authority across a domain, the migration criterion becomes difficult to satisfy. Specialist subpopulations that previously maintained distinctions outside the official scheme find their maintenance topologies disrupted by the same institutional forces that imposed the compression. Distinctions migrate toward Region I not because their reachability divergence has decreased but because the maintenance topologies that previously sustained them have been dismantled. The result is not merely that some distinctions disappear. It is that the *capacity to maintain* distinctions outside the dominant scheme degrades, narrowing the range of futures that remain navigable for agents who do not control the compression.

7. Conclusion

We began from a simple inversion: distinctions are not merely represented; they are maintained. A distinction survives because a population continues paying the cost of preserving it. This reframing — from informational objects to maintained objects — is not a metaphor. It is the primitive from which the framework's results follow.

Three contributions have organized the argument. First, distinction migration is maintenance reallocation. A distinction does not simply survive in a specialist community because it retains semantic value. It survives because a smaller population is willing to bear a larger per-capita maintenance cost. What changes during migration is the maintenance topology — the distribution of who pays to keep the distinction alive — not the semantic content of the distinction itself.

Second, repair pressure is driven by attributed failure rather than objective failure. The legibility operator $\mathcal{L}_\delta$ separates what goes wrong from what is recognized as going wrong. This separation explains why some collapsed distinctions generate rapid repair dynamics and others impose continuous objective costs without ever producing attributed pressure sufficient to enter the repair-seeking regime. Illegibility is not a delay; it is a structural condition.

Third, canonical restoration requires repair generation and repair coordination as independent conditions. A repair that is available — achieving high Q^* — may not be adopted because $\mathrm{Var}_\pi(Q)$ remains high. A population in which adoption pressure is high may not generate a canonical repair because $A(t)$ is low. The characteristic lag between paradigm availability and paradigm adoption is a consequence of this independence: the conditions for repair generation and repair coordination are satisfied at different times by different mechanisms.

On paradigm shift

The framework offers a reinterpretation of scientific revolutions as transitions in a three-variable dynamical system rather than as changes in belief or worldview. A paradigm shift is the visible phase transition that occurs when three conditions coincide: repair pressure rises above the maintenance threshold; a high-quality repair becomes available given the current adjacent possible; and coordination disagreement collapses below the adoption threshold. Each condition has its own dynamics. The historical duration of scientific revolutions is largely determined by how long it takes for the slowest of these three processes to complete.

This account preserves what Kuhn got right — the reality of theoretical turbulence, the lag between availability and adoption, the role of community structure in determining when transitions occur — while making each component precise and separable. The separation matters for intervention: a community in Region III-a needs a different response than one in Region III-b, and identifying which regime a given dispute inhabits is now a tractable question rather than a matter of interpretation.

Open questions

Several questions the framework raises remain open. The legibility operator $\mathcal{L}_\delta$ is characterized qualitatively here; a formal treatment would require specifying how attribution is established and how its delay is determined by failure-mode structure. The adjacent-possible availability term $A(t)$ is treated as largely exogenous to distinction dynamics; a more complete theory would explain what raises and lowers $A(t)$ endogenously, including the feedback effects between repair-seeking effort and adjacent-possible availability. The coordination coefficient κ is identified with institutional and demographic structures but its quantitative determinants remain unspecified. And the relationship between the framework developed here and formal models of belief revision, social choice, and institutional change is unexplored territory.

Recent work in cognitive science has converged on a closely related question from the opposite direction: whether communicative need is sufficient to explain which distinctions languages maintain [13]. The maintenance-topology framework suggests that communicative need is necessary but not sufficient — it determines R_D^θ and C, but whether those quantities generate sufficient repair pressure to sustain a distinction depends also on M, S, and the distribution of maintenance costs across the relevant population. Connecting the two approaches formally is a productive direction for future work.

What the framework offers is a common vocabulary for phenomena that have been studied separately: the survival of technical distinctions in specialist communities, the failure of locally successful interventions to achieve canonical status, the long lag between theoretical availability and theoretical adoption, the illegible accumulation of failure costs in bureaucratic and machine-learning systems, and the political economy of who bears the cost of imposed compression. These phenomena have a common structure. The lifecycle of distinctions is governed by maintenance costs, attribution dynamics, and coordination mechanisms — and the study of that lifecycle is, at its core, the study of how representational systems manage the tension between compression and navigability.

What the framework does not claim

The framework developed here is not a theory of intellectual progress. It does not predict that true or useful distinctions will eventually be recognized and adopted. It does not claim that canonical restoration is the typical outcome for collapsed distinctions, or that Region IV represents a stable endpoint rather than a temporary equilibrium. These would be Whiggish readings that the framework's own structure rules out.

What the framework actually entails is considerably less comfortable.

A distinction can be true and dead. Truth is not a maintenance-topology. A distinction that tracks real features of the world will collapse if no population is willing to bear its maintenance cost, and will remain collapsed regardless of its accuracy.

A distinction can be false and alive. If a distinction serves the interests of the populations maintaining it, it can persist in Region II or Region IV indefinitely. The framework explains persistence but does not certify the validity of what persists.

A distinction can be alive in one population and dead in another simultaneously. The automation-augmentation distinction illustrates this: it is in Region III-b for researchers who study human-machine complementarity and in Region I for accounting regimes that make augmentation value illegible. Neither state is the "true" state of the distinction. Both are real, and they are real for different populations.

A distinction can disappear from the dominant representation while remaining continuously maintained in a specialist topology. This is Region II. The distinction is not lost. It has changed who pays for it. Nothing in the framework implies that specialist maintenance is a transitional state on the way to restoration. It may be the permanent equilibrium.

Canonical restoration is a rare outcome produced by the coincidence of three independent conditions. Most collapsed distinctions do not restore. They either die, migrate, or cycle through turbulence without resolution. The historical cases that achieve Region IV are precisely the ones visible in retrospect because they succeeded — but the framework predicts that for every plate tectonics, there are many distinctions that entered turbulence and never exited, whose absence from the intellectual record is itself a product of the dynamics described here.

The framework is therefore not optimistic about the fate of distinctions. It is structural. It describes the conditions under which distinctions survive, migrate, and die, without assuming that any of these outcomes is more natural or more final than the others. The appropriate attitude it recommends is not confidence that good distinctions will eventually prevail, but attention to the maintenance topologies currently in place and the populations currently bearing their costs.

8. Future Directions

The framework developed here treats each distinction-domain-population triple (D, θ, Π) as an independent dynamical unit. Four natural extensions would relax this independence and connect the maintenance-topology approach to broader questions in institutional theory, network science, and reachability geometry.

Coupled population dynamics

The paper establishes that distinct populations may occupy different regions of the phase space simultaneously for the same distinction. The populations are nonetheless treated as independent: repair pressure in one does not directly influence repair pressure in another. A natural extension is a coupled dynamical system:

$$\frac{dP_r^{(\Pi_i)}}{dt} = \alpha_i E_a^{(\Pi_i)} - \beta_i S^{(\Pi_i)} + \sum_j \gamma_{ij} \, P_r^{(\Pi_j)}, \tag{9}$$

where γ_{ij} measures the influence of population Π_j's repair pressure on population Π_i's. Positive γ_{ij} would model the export of repair urgency from specialist communities — whose attributed failures are legible — outward to general populations where the same distinction is in Region I. Negative γ_{ij} would model competitive dynamics where one population's adoption of a workaround reduces another's attributed pressure. This extension would allow the theory to describe how distinctions move between maintenance regimes across population boundaries rather than only within them.

Distinction ecologies and maintenance budgets

The current framework considers one distinction at a time. In practice, distinctions compete for shared maintenance resources: attention, institutional overhead, curricular space, and cognitive capacity are all finite. A distinction ecology would introduce a maintenance budget constraint:

$$\sum_i M(D_i, \theta, t, \rho) \ \leq \ B(\theta, t), \tag{10}$$

where $B(\theta, t)$ is the total maintenance capacity available to domain θ at time t. Under this constraint, the survival of one distinction can come at the cost of another. Institutions that expand their operational scope without expanding their maintenance capacity would predictably shed distinctions even when individual repair pressures remain high. This connects the maintenance-topology framework to questions in curriculum design, institutional overload, and the representational bottlenecks that arise in large-scale machine-learning systems.

Network topologies of maintenance populations

The paper treats maintenance populations abstractly, characterizing them by aggregate quantities such as $M_{\text{specialist}}$ and κ. A more refined treatment would model the population as a network $G = (V, E)$, where nodes are individual maintainers and edges represent coordination channels. Migration stability (Proposition 3.6) and the coordination bottleneck (Proposition 4.3) would then depend not only on aggregate κ but on the connectivity structure of G. Sparse or fragmented networks might lose distinctions even when the per-capita maintenance cost is theoretically affordable, because the coordination channels required to sustain collective maintenance are absent. This direction connects the framework to social network analysis and to the literature on the maintenance of common-pool resources [7].

Reachability fields and the geometry of distinction maintenance

The repair quality measure already uses recovered admissible future volume as its numerator:

$$Q(R \mid \pi) \ = \ \frac{\Delta\,\mathrm{Vol}(\mathcal{A})}{\Delta K(R \mid \pi)}.$$

A natural further step is to define a reachability field over the space of distinctions:

$$\Phi_D \;=\; \mathrm{Vol}(\mathcal{A}_D), \tag{11}$$

where $\mathcal{A}_D$ is the admissible future set unlocked by maintaining distinction D. Distinction maintenance then becomes a special case of reachability preservation: the problem of maintaining a distinction is the problem of preserving a region of the admissible future manifold against compression. This connects the present framework to a broader program in which reachability geometry serves as the unifying object across physics, cognition, computation, and institutional structure. In that setting, the current paper's results — migration, hysteresis, the legibility bound, the coordination bottleneck — would appear as theorems about how compression schemes affect the topology of admissible future spaces, and canonical restoration would be interpretable as the recovery of a connected component of $\mathcal{A}$ that had been severed by the collapse of D.

The author has no institutional affiliation to acknowledge. All errors are the author's own.

References

[1] Kuhn, T. S. *The Structure of Scientific Revolutions*. University of Chicago Press, 1962.

[2] Lakatos, I. *The Methodology of Scientific Research Programmes*. Cambridge University Press, 1978.

[3] François, A. "Semantic maps and the typology of colexification." In *From Polysemy to Semantic Change*, ed. M. Vanhove. John Benjamins, 2008.

[4] Xu, Y., Regier, T., and Malt, B. C. "Historical semantic chaining and efficient communication: The case of container names." *Cognitive Science*, 40(8):2081–2094, 2016.

[5] Brochhagen, T., Boleda, G., Gualdoni, E., and Xu, Y. "From language development to language evolution: A unified view of human lexical creativity." *Science*, 376(6598):1210–1215, 2022.

[6] Holling, C. S. "Resilience and stability of ecological systems." *Annual Review of Ecology and Systematics*, 4:1–23, 1973.

[7] Ostrom, E. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990.

[8] March, J. G. "Exploration and exploitation in organizational learning." *Organization Science*, 2(1):71–87, 1991.

[9] Kauffman, S. *Investigations*. Oxford University Press, 2000.

[10] Johnson, S. *Where Good Ideas Come From: The Natural History of Innovation*. Riverhead Books, 2010.

[11] Jackson, S. J. "Rethinking repair." In *Media Technologies: Essays on Communication, Materiality, and Society*, ed. T. Gillespie, P. J. Boczkowski, and K. A. Foot. MIT Press, 2014.

[12] Star, S. L. "The ethnography of infrastructure." *American Behavioral Scientist*, 43(3):377–391, 1999.

[13] Beekhuizen, B. "Commentary on lexical distinction maintenance and communicative need." *Nature Human Behaviour*, 2024.