

# The Admissibility Crisis

On the Collapse of Future Option Space

**Flyxion**

June 27, 2026

## Abstract

This paper formalizes the structural dynamics of the metacrisis by pivoting from standard resource-depletion models to a rigorous geometric framework of state-space reachability and structural coherence ( $\mathcal{R}_A(t)$ ). By synthesizing Ellul's *technique*, Baudrillard's *hyperreality*, Doctorow's *chokepoints*, and Mattei's *austerity*, we map how complex socio-technical systems actively optimize for their own status quo through the systematic foreclosure of future states. We introduce the Carelessness Principle ( $\frac{\partial M}{\partial x} \approx 0$ ) to model systemic externalization without requiring individual malice, and detail how non-linear feedback loops—the sigmoidal repair floor  $\phi(\mathcal{V}_A)$  and the hyperbolic distinction loss term  $k/\mathcal{V}_A$ —trigger sudden bifurcation cascades at critical volume thresholds ( $\mathcal{V}_{\text{crit}}$ ). We further introduce a novelty generation term  $N(t)$  to distinguish repair of prior admissible states from the boundary reconfiguration through which genuinely new regions of state space become accessible. The framework yields a coupled nonlinear dynamical system whose variables are empirically proxiable, whose central claim—that  $\frac{d\mathcal{V}_A}{dt} < 0$  under current conditions—is falsifiable in principle, and whose core inversion—that admissibility is prior to preference, not downstream of it—distinguishes this account from welfare-economic and game-theoretic approaches to civilizational risk.

## I. The Wrong Diagnosis

The standard account of civilizational crisis is a story about scarcity. We are running out of clean water, arable land, functioning institutions, shared truth, civic trust, attention, time. The crisis, on this account, is a depletion problem. We had enough of something; now we are running short; eventually we will run out.

This framing is wrong—not merely incomplete, but wrong in a precise sense that matters for what follows. The problem is not that we lack inputs. The problem is that we are systematically destroying the conditions under which inputs remain usable. A civilization running out of *admissible futures* looks identical, from the inside, to one running out of resources. The symptoms are the same. The diagnosis is different. The treatment is not the same at all.

Consider what it would mean for a system to be in genuine resource crisis versus admissibility crisis. A village running out of water has a depletion problem: the aquifer drops, the wells go dry, eventually there is nothing left to pump. A village whose water infrastructure has been systematically optimized for throughput at the expense of maintenance, whose engineers have been replaced by quarterly reporting cycles, whose pipes have been privatized to firms whose metrics do not include long-term structural integrity—that village may have a full aquifer and still find itself without water. The resource is present. The capacity to reach it has been destroyed.

The metacrisis—that cluster of interlocking civilizational failures that commentators increasingly struggle to name as a single thing—is of the second kind. It is not primarily a story about running out. It is a story about the progressive contraction of the set of futures we can still reach.

This reframing has a precise consequence for how we understand contemporary power. The defining problem of the modern economy is not concentration of ownership but concentration of reachability. The most dangerous forms of power are not those that seize assets but those that govern the conditions under which assets remain usable—the pathways through which activity must pass, the metrics by which value is recognized, the distinctions through which reality is organized. Everything that follows in this essay is a consequence of that claim.

## II. Admissibility Defined

The notion of admissibility, in its technical sense, concerns which states a system can coherently occupy. A state is admissible if it can be reached from the current configuration by a sequence of operations that preserve the system’s internal coherence—not optimality, not efficiency, not even desirability. Just coherence: the capacity of the system to continue functioning as the kind of system it is.

Admissibility is prior to preference. You cannot optimize your way toward an inadmissible state; you can only be surprised by it when you arrive. A chess position can be legal without being desirable; it can be desirable without being reachable from where you currently stand. The set of reachable positions—the option space—is a structural fact about the game, not a matter of preference or strategy. The strategy operates within the option space. It does not define it.

We formalize this as follows. Model the configuration of a civilization at time  $t$  as a point  $x(t)$  on a smooth high-dimensional manifold  $\mathcal{M}$ —the space of all possible social, institutional, ecological, and material configurations. The *distinction infrastructure* of the civilization—its categories, roles, relationships, meanings, institutional forms, and shared knowledge—imposes a set of functional constraints on which states remain coherently occupiable. Represent these constraints as a time-varying boundary operator  $B(t)$ : a state  $x \in \mathcal{M}$  is *admissible* at time  $t$  if and only if

$$B(t)x \geq 0.$$

The constraints encoded in  $B(t)$  are not preferences or policies. They are structural conditions: the minimal requirements for the system to continue functioning as the kind of system it is. A democracy requires electoral infrastructure, shared epistemic standards, and trust in institutional continuity. A language requires speakers and transmission mechanisms. A scientific community requires shared methodological norms and accessible archives. When  $B(t)$  degrades—when the constraints it encodes collapse—the set of admissible states contracts.

The *admissibility volume* at time  $t$  is then the measure of the set of states that satisfy the constraints at that time:

$$\mathcal{V}_A(t) = \mu\left(\{x \in \mathcal{M} \mid B(t)x \geq 0\}\right),$$

where  $\mu$  is a measure on the state manifold  $\mathcal{M}$ . This is the set of states that are structurally coherent at time  $t$ . It is not identical to the set of states that are reachable. Admissibility and reachability are distinct: a state can be admissible but unreachable from the current position, just as a chess position can be legal without being reachable from the current configuration through any sequence of legal moves.

Reachability requires dynamics. Define the reachable set at time  $t$  as

$$\mathcal{R}(t) = \{x' \in \mathcal{M} \mid x' = \Phi_t(x_0, u)\},$$

where  $\Phi_t(x_0, u)$  is the flow of the system from initial state  $x_0$  under available interventions  $u(\cdot)$ . The object that governs what a civilization can actually do is the intersection of the admissible and

the reachable:

$$\mathcal{R}_A(t) = \mathcal{R}(t) \cap \mathcal{V}_A(t).$$

We call  $\mathcal{R}_A(t)$  the *admissible reachable set*. Contraction of  $\mathcal{R}_A(t)$  can occur through two independent channels:  $\mathcal{V}_A(t)$  shrinks as  $B(t)$  degrades, making previously coherent states structurally unavailable; or  $\mathcal{R}(t)$  shrinks as path dependence, chokepoint concentration, and political constraint remove trajectories from the reachable set without necessarily making the destination states inadmissible. Both channels are active in the mechanisms described in this essay. The rate  $\frac{d\mathcal{V}_A}{dt}$  approximates the first channel; chokepoint capitalism and austerity governance are the primary drivers of the second.

When  $B(t)$  degrades faster than repair interventions  $u(\cdot)$  can compensate,  $\mathcal{V}_A(t)$  contracts and  $\mathcal{R}_A(t)$  contracts with it. When  $B(t)$  is actively maintained and  $\mathcal{R}(t)$  is kept redundant through distributed pathways,  $\mathcal{R}_A(t)$  expands.

This changes admissibility from an intuition into a structural quantity. A democracy does not merely become “weaker”; it approaches a boundary in  $\mathcal{M}$  below which democratic repair lies outside  $\mathcal{R}_A(t)$ —neither reachable from the current state nor structurally coherent given current constraints. A language does not merely “decline”; it crosses a threshold in  $B(t)$  below which recovery lies in  $\mathcal{V}_A(t) \setminus \mathcal{R}(t)$ : coherent in principle, but unreachable without discontinuous external intervention. The admissibility crisis is therefore a measurable structural claim:  $\mathcal{R}_A(t)$  is currently shrinking faster than repair activity can compensate.

The central claim of this essay is that the metacrisis represents a contraction of  $\mathcal{R}_A(t)$ , driven primarily by degradation of  $B(t)$  and concentration of  $\mathcal{R}(t)$ , and that the mechanisms driving that contraction are not independent failures but expressions of a single underlying dynamic: a systematic bias, built into the optimization structures of modern technological capitalism, against the activities that maintain both.

### III. Local Rationality, Global Inadmissibility

Before developing the formal account, it is worth dwelling on what makes the admissibility crisis so difficult to diagnose. Each of its generating mechanisms is, in isolation, locally rational. This is what distinguishes it from simple dysfunction.

Engagement maximization is rational for a platform facing quarterly earnings review. The algorithm that keeps users on-screen longer generates more advertising revenue. The engineers who build it are doing their jobs. The executives who reward them are doing theirs.

Surveillance advertising is rational for a firm that needs signal. Behavioral data produces better-targeted ads. Better-targeted ads convert at higher rates. Higher conversion rates justify the data infrastructure.

Algorithmic labor management is rational for an operator minimizing variance. Human managers introduce inconsistency. Algorithmic systems enforce standardization. Standardized labor processes are easier to scale.

Planned obsolescence is rational for a manufacturer who needs demand. A product designed to last twenty years will not be replaced for twenty years. A product designed to last three years will generate seven replacement cycles over the same period.

Financialization is rational for capital seeking yield. Patient investment in productive capacity requires long time horizons and accepts interim underperformance. Financial instruments indexed to near-term earnings metrics produce more predictable returns.

None of these are irrational. Each follows from premises that, within the relevant optimization context, are entirely sound. And yet collectively they constitute something that no individual optimizer intended and no individual decision produced: a systematic tax on admissibility. Not on any particular outcome. On the option space itself.

This is the key structural move. The harm produced by these mechanisms is not harm to any specific future. It is harm to the reachability of futures in general. They do not make any particular outcome worse. They shrink the set of outcomes we can still reach.

[O

#### **IV. The Compression Trap**

To understand how local optimization produces global admissibility collapse, we need to understand what happens to the representations that optimization systems use to navigate.

Jacques Ellul, writing in the mid-twentieth century, identified a dynamic he called *la technique*: the tendency of modern societies to reorganize themselves around the pursuit of maximum efficiency through rationalized procedures. His concern is often misread as anti-technology. His actual argument is more precise and more disturbing: that once efficiency becomes the dominant criterion, institutions begin selecting for what can be optimized, and progressively eliminating what cannot.

In admissibility terms, technique introduces a selective pressure on future states. Before optimization, the admissibility volume contains many possible futures—some efficient, some inefficient, some beautiful, some ugly, some durable, some fragile. After sustained optimization pressure, the

admissibility volume contracts toward the subset of futures compatible with technical efficiency:

$$\mathcal{V}_A \longrightarrow \mathcal{V}'_A = \{ s \in \mathcal{V}_A \mid s \text{ is legible to optimization} \}.$$

The system gains local performance while losing global optionality.

This is why technically optimized systems so often exhibit a characteristic combination of power and brittleness. A hospital optimized for throughput becomes excellent at handling predictable cases and fragile in the face of unusual ones. A supply chain optimized for cost efficiency performs extraordinarily well under stable conditions and catastrophically under disruption. A university optimized for measurable research metrics becomes very good at producing citable papers and increasingly poor at producing intellectual novelty. The optimization worked. The problem is that success itself became a contraction operator on admissible futures.

Ellul's deeper insight—often underemphasized—was not merely efficiency but *autonomy*. Technique gradually becomes self-reproducing because every technical solution generates secondary problems that require further technical solutions. The system begins expanding its own necessity. In admissibility terms:

$$T_{n+1} = f(T_n),$$

where each technical intervention produces a new state whose internal coherence requires additional technical intervention to maintain. Optimization landscapes become self-expanding. Consider the trajectory a major social platform is embedded within: engagement optimization creates attention fragmentation; attention fragmentation requires recommendation systems; recommendation systems produce filter bubbles; filter bubbles require content moderation; content moderation requires machine learning at scale; machine learning requires surveillance-scale data collection; data collection creates privacy problems; privacy problems require new governance technologies. Each solution is technically rational. Each one expands the optimization landscape without closing the underlying admissibility deficit.

What Ellul observed at the level of institutions, Jean Baudrillard observed at the level of representation. His central concern—often sensationalized, less often understood—was that representations cease functioning as representations. The map becomes more operationally important than the territory. The sign detaches from what it once signified and begins to determine it.

Consider what happens when GDP becomes the primary instrument by which governments navigate economic reality. Initially GDP is a useful compression: it summarizes a vast array of economic activity into a single number that is tractable enough to govern with. Over time institutions organize around it. Policies are evaluated by their effect on GDP. Careers are built on GDP

growth. Elections are won and lost on GDP figures. Eventually the representation acquires more political and economic force than the system it was originally designed to describe.

At this point something structurally important has happened. The projection  $\pi(X)$ —the compressed representation of the economic system  $X$ —has begun determining  $X$  rather than merely describing it. The causal arrow has reversed. But Baudrillard’s strongest formulation goes further than simple reversal. Once a representation becomes causally dominant, the system itself reorganizes around it:

$$X' = F(X, \pi(X)).$$

The representation is no longer a description of  $X$ ; it is one of the forces *constituting*  $X'$ . A company’s stock price affects hiring decisions, borrowing costs, acquisition strategy, executive compensation, media coverage, and employee morale. At some point the stock price is not merely a representation of the company. It is a causal mechanism constituting what the company is. The map has entered the territory. The territory reorganizes around the map.

We can make this formally precise. Define the true state of the social system as  $S(x)$  and the optimization metric as  $M(x) = \pi(S(x))$ . Under competitive market pressure, institutions navigate not along the gradient of  $S$  but along the gradient of  $M$ . The actual trajectory of the system is governed by a reorganization vector field:

$$\dot{x} = \nabla M(x) + \eta(x),$$

where  $\eta(x)$  represents autonomous human agency and residual non-optimizing dynamics. When  $\pi$  is a faithful compression,  $\nabla M \approx \nabla S$  and the trajectory tracks the underlying social reality. As divergence grows, however, the directional derivative of the map increasingly departs from the directional derivative of the territory.

Because  $S(x)$  is not directly observable, the compression distortion term  $C(t)$  cannot be computed as a direct divergence between  $\nabla M$  and  $\nabla S$ . A more tractable formulation defines  $C(t)$  in terms of the observable *reorganization velocity*—the rate at which institutions and social systems are reorienting themselves around metric proxies rather than substantive outcomes. Formally:

$$C(t) = \gamma \left\| \frac{d}{dt} [X' - F(X, \pi(X))] \right\|,$$

where  $\gamma > 0$  is a coupling constant and the norm measures how rapidly the system is reorganizing around its own representations. This quantity is empirically tractable: it can be approximated by tracking institutional drift toward proxy optimization (the rate at which universities restructure around citation metrics, the rate at which healthcare systems restructure around billing codes, the

rate at which media organizations restructure around engagement rates). When institutions reorganize rapidly around metric proxies,  $C(t)$  is large. When they maintain substantive governance of underlying systems,  $C(t)$  is small. The Compression Trap is not a metaphor. It is a condition on the reorganization velocity—a condition that intensifies monotonically under competitive optimization pressure.

This is why metric capture is so difficult to reverse. Once  $\pi(X)$  becomes constitutive of  $X'$ , removing it changes the system itself. You cannot simply return to measuring what you originally cared about, because the system being measured has been reshaped by the measurement.

The examples multiply easily: engagement replacing genuine communication; citation counts replacing understanding; follower counts replacing actual expertise; credit scores replacing trustworthiness; test scores replacing education; stock price replacing productive capacity. Each begins as a useful approximation. Each becomes, under optimization pressure, a hyperreal object—more causally real than what it was meant to describe.

The Compression Trap, then, is this: technique requires compressed metrics; compressed metrics become optimization targets; targeted metrics reorganize the systems they were meant to describe; institutions begin navigating by maps that no longer track the territory—because the territory has been partly remade in the map’s image.

Ellul and Baudrillard are diagnosing complementary stages of the same process. Ellul explains why optimization becomes dominant and self-reproducing. Baudrillard explains what happens after optimization dominates: the representation becomes constitutive. Together they describe the engine and the pathology of admissibility compression.

## V. The Optimization Gradient

Before proceeding to repair, it is worth making the structure of the problem precise, because precision here has a significant consequence: it renders the question of individual moral failure largely irrelevant.

Let  $M(x)$  be a measurable proxy for some true social value, and let  $S(x)$  be the actual social state. Organizations embedded in markets face strong selection pressure to optimize  $M(x)$ , because  $M$  is what can be measured, reported, rewarded, and penalized at the timescales on which accountability operates. Define the divergence

$$D = \|M - S\|$$

as a structural quantity. As optimization pressure increases,  $D$  grows. The optimization gradient

points toward maxima of  $M$ , not maxima of  $S$ . The harder institutions optimize, the further trajectories drift from the underlying social reality.

This means the problem is not corruption in the ordinary sense. No executive needs to intend ecological destruction. No engineer needs to want social fragmentation. No advertiser needs to want the erosion of epistemic infrastructure. The gradient just points in those directions, and institutions follow gradients. That is what it means to be an institution embedded in a competitive market: you follow the gradient or you are replaced by something that does.

This also explains why replacing individual leaders rarely changes outcomes. A new executive in the same optimization landscape faces the same gradient. The attractors remain the same. The trajectories remain the same. The problem is not the people; it is the field.

Two concrete cases make the variables real and measurable.

**Materialism and the externalized substrate.** Consumer electronics provide a clean example. A smartphone manufacturer optimizes for margin, unit volume, and replacement rate. The metrics are clear. What they do not contain: the labor conditions in the factories where devices are assembled—often in jurisdictions with deliberately weakened labor protections that make low wages structurally available to the supply chain—the toxicity of the recycling processes in the countries where those devices are ultimately dismantled, the ecological costs of rare earth extraction, the systemic effects of designing products for obsolescence rather than repair. These costs are real. They accumulate in systems—bodies, ecosystems, labor markets, communities—that  $M(x)$  does not include.

The device is cheap because its true costs are externalized. The externalization is not an accident or a moral failing. It is a structural consequence of optimizing  $M$  when  $D = \|M - S\|$  is large. The metric measures what can be monetized within the relevant jurisdiction and accounting period. The substrate—the ecological, social, and material systems on which production depends—lies outside the measurement boundary.

The admissibility consequence: repair capacity for the externalized systems is never built, because it never appears in any optimization loop. Over time the substrate degrades while the metrics remain healthy, until the degradation becomes too large to contain and crosses back into  $M$ . At that point the system discovers it has been consuming admissibility volume for the duration of the optimization cycle.

**Industrial meat and the marketing of the compression.** The global meat industry offers a different and in some ways more instructive case. The industry faces a genuine problem: the

system it operates—factory farming at scale—generates costs that would, if widely understood, reduce demand. The admissibility costs are substantial: antibiotic resistance accumulating in the global microbiome, emissions accumulating in the climate system, water systems stressed by runoff, labor conditions in processing facilities that would not survive public visibility, ecological systems stripped to produce feed grain.

The industry’s response has been to spend at massive scale on the repair of the *representation* rather than the repair of the *reality*. Billion-dollar marketing campaigns do not improve the underlying system. They repair the gap between  $M$  and  $S$  by adjusting  $M$ —by making meat appear normal, natural, healthy, affordable, and culturally indispensable—rather than by reducing  $D$  through changes to  $S$ .

This is the most important structural feature of the case: genuine repair capacity exists. Resources, attention, and organizational competence are being deployed on repair. But the repair is directed at the compression rather than the substrate. The map is being maintained while the territory deteriorates. The industry is not an outlier. It is the general form that repair takes when the optimization gradient points away from the underlying system.

## VI. The Repair Deficit

Every durable system continuously repairs the distinctions it depends upon. This is not a metaphor. It is a structural requirement.

A democracy depends on distinctions: between legitimate and illegitimate authority, between public and private interest, between evidence and fabrication, between elected and appointed, between law and preference. These distinctions do not maintain themselves. They require continuous active repair—through institutions, through journalism, through civic participation, through education, through the accumulated practices of political culture. When repair falls below a threshold, the distinctions begin to blur. When they blur sufficiently, the system crosses into inadmissibility: the state space of functional democratic governance becomes unreachable without discontinuous intervention.

The same holds for science, language, community, ecological systems, and material infrastructure. In each case the system is not a static object that degrades through use. It is an ongoing process of distinction repair that produces, as its output, the appearance of a stable system. The stability is real; it is just not free.

Platform firms are structurally positioned to capture the value produced by repaired systems while externalizing the repair cost. Social media monetizes attention that is only worth monetizing be-

cause shared language, social norms, and epistemic infrastructure have been maintained by other means. Search engines monetize knowledge that is only worth indexing because scientific institutions, universities, journalism, and libraries have spent centuries building and repairing it. Cloud platforms monetize communities that formed and developed on open infrastructure that the platforms did not build and do not maintain.

This is not necessarily malicious. It is structurally invisible. Repair does not appear on the balance sheet. The distinction infrastructure that makes a platform's product valuable is a commons that long predates the platform; its degradation is slow; its cost is distributed across the whole system; and its contribution to any particular metric is immeasurable. So it does not get measured. So it does not get repaired.

The result is a civilization-scale repair deficit. The substrate of social, epistemic, ecological, and material distinctions is being consumed faster than it is restored. Each extraction is locally rational. The aggregate is a progressive narrowing of the admissibility volume.

## VII. Platforms as Distinction Governors

A natural objection to the foregoing is that these mechanisms predate digital platforms. Planned obsolescence existed before the smartphone. Metric capture preceded social media. The repair deficit long preceded cloud computing. If the mechanisms are not new, what is?

The answer is not that platforms created the problem. The answer is that digital systems dramatically accelerate admissibility contraction because they operate directly on distinction infrastructure—on the mechanisms by which distinctions themselves are produced, propagated, reinforced, forgotten, and repaired.

A steel company affects physical infrastructure. A social media platform affects the ecology of distinctions. This is a difference in kind, not merely degree. The platform is not merely another corporation; it is an institution that partially governs the process by which shared reality is constructed and maintained.

Historically, distinction production was distributed across many partially independent institutions. Libraries preserved what was worth knowing. Newspapers established what was publicly visible. Universities determined what counted as knowledge. Local communities maintained tacit expertise. Courts adjudicated what was legitimate. Scientific societies repaired consensus. Churches maintained moral distinctions across generations. None of these institutions was neutral; all of them were imperfect; many of them failed badly in various ways. But they were *distributed*. No single optimization gradient governed all of them simultaneously.

Platforms centralize increasing fractions of this process into a handful of optimization systems. Search rankings determine discoverability. Recommendation algorithms determine visibility. Moderation systems determine legitimacy. Large language models increasingly determine reconstructability—what can be recovered from partial information. The ecology of distinctions that was once governed by many partially independent institutions, each with different incentive structures and failure modes, is increasingly governed by a small number of systems whose gradients all point in the same direction.

The danger is not censorship, bias, or misinformation in isolation, though all of these are real. The danger is that the civilization’s distinction repair machinery has become coupled to a small set of metrics whose optimization gradients are structurally unrelated to long-term admissibility. When a library fails, the distinctions it maintained degrade locally. When the search ranking system that has replaced it fails—or, more precisely, succeeds at optimizing its metric—the degradation is global and simultaneous.

This is why the Big Tech critique, when framed as a matter of corporate power or content moderation failure, consistently misses its own target. The problem is not that these institutions are powerful. It is that their power is exercised over the substrate from which all other institutions draw their coherence, and the optimization gradients governing that exercise have no term for admissibility.

## **VIII. Chokepoint Capitalism**

Not all forms of power operate by ownership. Some operate by controlling passage.

Historically, chokepoints were physical. Mountain passes, ports, canals, bridges, and rail junctions concentrated flows into narrow regions where they could be monitored, taxed, delayed, or redirected. The strategic significance of these locations arose not from what they contained but from what moved through them. Whoever controlled the passage controlled the surrounding territory indirectly.

Contemporary capitalism increasingly operates through an analogous logic. The most powerful firms often do not produce the goods being exchanged, create the content being consumed, write the software being distributed, or generate the knowledge being indexed. Instead they occupy positions through which these activities must pass. Their power derives from the ability to govern access rather than from direct production.

App stores govern software distribution. Search engines govern discoverability. Payment processors govern transactions. Cloud providers govern computation. Social platforms govern visibility. Ad-

vertising networks govern attention. Recommendation systems govern exposure. The most valuable position in a modern economy is often not participation in a flow but control over the conditions under which the flow occurs.

This distinction matters because chokepoints alter the geometry of admissibility.

In a distributed system, many paths connect a present state to a future state. If one route fails, alternatives remain available. Repair can proceed through redundancy. Adaptation is possible because multiple trajectories remain open. The admissibility volume remains large not because every path is efficient but because many paths exist.

A chokepoint changes this structure. Numerous trajectories collapse into a small number of privileged channels. The apparent efficiency of the system increases. Coordination costs fall. Friction decreases. The network becomes easier to navigate. Yet simultaneously the number of independent routes declines. What appears as optimization from the perspective of throughput appears as contraction from the perspective of reachability.

The danger is not monopoly in the classical sense. Traditional monopoly concerns ownership and pricing power. Chokepoint power concerns path dependence. A firm may allow competitors, users, creators, and communities to operate freely while still retaining decisive influence over the conditions of their operation. The relevant question is not who owns the activity but who governs the pathways through which the activity becomes possible.

This produces a characteristic asymmetry. Participants remain formally autonomous while becoming structurally dependent. A journalist may be free to publish, yet dependent on algorithmic visibility. A software developer may be free to create, yet dependent on platform approval. A business may be free to sell, yet dependent on payment infrastructure. The appearance of decentralization coexists with increasing concentration of effective control.

The admissibility consequences extend beyond economics. Chokepoints transform failures that would otherwise remain local into failures with system-wide effects. A degraded search index affects the discoverability of knowledge across entire domains. A change in recommendation algorithms alters the visibility landscape for millions simultaneously. A failure in cloud infrastructure propagates through countless dependent systems. The concentration of pathways reduces the capacity of the larger system to absorb shocks through rerouting.

More importantly, chokepoints distort repair incentives. Repair activity occurring outside the chokepoint becomes economically invisible. Communities maintain knowledge commons, volunteers moderate discussions, researchers produce information, and local institutions preserve expertise. Yet value accrues disproportionately to the infrastructure through which these activities pass. The

result is another version of the repair deficit: the pathways capture increasing fractions of the value generated by the territory while contributing comparatively little to the regeneration of the systems that make the pathways worth traversing.

This dynamic explains why contemporary capitalism often appears simultaneously innovative and fragile. Innovation continues because optimization within the pathways remains highly effective. Fragility increases because the pathways themselves become single points of admissibility failure. The system gains efficiency by sacrificing redundancy. It gains throughput by reducing optionality.

The central question is therefore not whether chokepoints are useful. Many are extraordinarily useful. The question is whether the efficiencies they generate exceed the admissibility volume they consume. A civilization whose future depends upon a small number of gateways may continue functioning normally for long periods. The cost appears only when repair becomes necessary and alternative routes no longer exist.

From the perspective of admissibility, the defining feature of chokepoint capitalism is not concentration of wealth. It is concentration of reachability. The system's future becomes increasingly dependent on a narrowing set of pathways through which all continuation must pass. The result is a civilization whose apparent freedom of movement conceals a progressive contraction of the routes by which movement remains possible.

## IX. The Governance of Admissibility

The mechanisms described so far share a common feature: admissibility contraction emerges from optimization systems pursuing their own internal logic. No one intends to shrink the future. The gradient just points that way. But this framing, while accurate for many mechanisms, omits a category that is both historically important and theoretically distinct: the deliberate governance of admissibility.

Political institutions do not merely allocate resources. They actively reshape the set of futures that remain reachable. The state is not a neutral arbiter operating within a fixed option space. It is itself an actor capable of expanding or contracting  $\mathcal{V}_A$  through policy, and the direction of that action is not predetermined.

Clara Mattei's historical analysis of austerity reveals this mechanism with unusual clarity. Her central argument is that the economic policies grouped under the term austerity—budget cuts, wage suppression, monetary tightening, fiscal discipline—are typically presented as technical necessities imposed by objective economic constraints. The historical evidence, she argues, supports a different reading. In the aftermath of the First World War, when democratic pressures were threatening to

expand the range of politically reachable economic arrangements, austerity policies functioned to restore social hierarchies by narrowing the option space itself.

In admissibility terms, this is not primarily a story about economics. It is a story about the use of governance to manage futures. Austerity reduces not merely expenditures but reachability. Labor movements become harder to sustain when the fiscal resources that support them are removed. Public ownership becomes harder to finance when capital markets are prioritized over public balance sheets. Social programs become harder to expand when deficit ceilings become constitutional constraints. Collective bargaining weakens when unemployment is maintained at levels sufficient to discipline wage demands. The state does not simply choose one future over another. It removes large regions of future possibility from the reachable state space:

$$\mathcal{V}_A \longrightarrow \mathcal{V}'_A \subsetneq \mathcal{V}_A.$$

The contraction itself becomes policy.

This mechanism is formally distinct from the others. Ellul's technique compresses admissibility through optimization pressure. Baudrillard's simulation distorts admissibility through representational capture. Chokepoint capitalism concentrates admissibility through pathway control. Mattei's austerity restricts admissibility through deliberate institutional constraint. The four mechanisms can operate simultaneously and reinforce each other. A government that has adopted austerity as a governing criterion will be less capable of funding the distributed distinction infrastructure that resists chokepoint capture, less capable of maintaining the repair capacity that resists the repair deficit, and less capable of sustaining the redundancy that resists reachability concentration.

Mattei also deepens the repair deficit analysis. Repair requires slack. It requires spare capacity and resources not immediately productive. A bridge cannot be maintained if maintenance budgets are continually cut to meet fiscal targets. A public health system cannot preserve resilience if every inefficiency is removed by austerity-driven restructuring. A university cannot preserve institutional memory if every department must justify itself according to short-term productivity metrics. One of Mattei's recurring historical observations is that austerity tends to destroy precisely the capacities that make long-term stability possible. In the formal notation:

$$R(t) \downarrow \quad \text{while policymakers report efficiency gains.}$$

The system appears healthier in the short run because visible expenditures decline. Its repair capacity is quietly consumed. This connects directly to Ellul: austerity becomes another instance of technique, efficiency gains achieved through contraction of admissibility volume, with the costs

appearing only when the system requires repair and discovers the capacity is no longer there.

Chokepoint capitalism and austerity governance produce complementary effects worth holding together. Chokepoints govern movement—they control the routes through which activity must pass. Austerity governs alternatives—it removes the institutional and fiscal conditions under which alternative routes could be built or maintained. One constrains paths. The other constrains destinations. A civilization may retain nominal freedom of action while losing practical reachability, because the paths are governed by private chokepoints and the destinations are governed by political constraint. The question that emerges from this combination is neither a market question nor a governance question in isolation. It is an admissibility question: which futures remain institutionally reachable, by whom, and through what remaining routes?

## **X. The Political Economy of Forgetting**

The memory architecture of modern technological capitalism is precisely inverted from what durable systems require.

Advertising infrastructure preserves every behavioral trace with extraordinary fidelity. A click made years ago persists in a targeting profile. An engagement pattern from a previous decade shapes what content a recommendation system surfaces today. Surveillance systems preserve movement, association, communication, and consumption across timescales and jurisdictions that human memory cannot approach.

Meanwhile: institutions forget long-term obligations. Governments forget policy failures from previous administrations. Communities forget accumulated expertise when practitioners retire without successors. Public discourse forgets context within a news cycle. Scientific fields forget negative results. Engineering institutions forget the reasoning behind safety constraints and then rediscover the hard way why those constraints existed. [I The problem is not memory scarcity. There is more recorded information today than at any previous point in human history. The problem is that the memory architecture is optimized for yield rather than coherence. What gets remembered is what generates monetizable signal. What gets forgotten is what generates the substrate within which signal can mean anything.

Ellul noticed a version of this: that technical systems preserve only information useful to technical administration. The records that survive are the records of optimization loops, not the records of what made those loops possible. Baudrillard noticed a complementary version: that the proliferation of information paradoxically destroys meaning, because meaning requires the maintenance of distinctions that information proliferation tends to blur. Each new piece of recorded information

that lacks contextual anchoring does not add to the total stock of distinction; it dilutes it.

The political economy of forgetting follows from the optimization gradient directly. If  $M(x)$  cannot capture the value of institutional memory, contextual knowledge, tacit expertise, or long-term obligation, then optimization pressure will systematically under-invest in these goods regardless of how much they contribute to  $S(x)$ . This is not a market failure in the ordinary sense. It is a case where the market mechanism is working exactly as designed, and the design does not include the goods that make admissibility possible.

This dynamic extends to how institutions respond to evidence of their own failures. An organization confronting evidence of divergence between its public representations and its operational realities faces a structural choice: repair the underlying system, or repair the representation of the system. Repairing the representation is typically cheaper, faster, and more legible to existing metrics. An organization can hire communications staff, issue clarifications, fund favorable research, or suppress inconvenient narratives far more efficiently than it can restructure the systems generating the failures those narratives describe. Optimization pressure therefore systematically favors representational repair over substrate repair. The result is a civilization increasingly capable of managing narratives about failure while becoming progressively less capable of addressing the failures themselves. What appears as public relations from the outside is often the visible manifestation of a deeper compression dynamic: the substitution of map maintenance for territory maintenance. The meat industry's marketing apparatus and the technology platform's moderation of criticism about itself are not different kinds of activity. They are the same activity at different scales.

## **XI. The Carelessness Principle**

The mechanisms described in previous sections—compression, simulation, repair deficits, platform capture, chokepoints, austerity governance, memory inversion—share a surface feature that is easy to misread. They look, from the outside, like failures of attention: like institutions that simply did not notice what they were doing to the systems they inhabited. But carelessness is not the same as inattention, and this distinction matters for how we understand what has gone wrong.

Large optimization systems produce a characteristic form of carelessness that is entirely compatible with high individual intelligence, careful internal analysis, and even genuine moral seriousness among the people working within them. The carelessness is not personal. It is structural. The system no longer needs to know what it depends upon, because what it depends upon does not appear in the variables it optimizes.

We can formalize this precisely. Define the *care* a system has for variable  $x$  as the sensitivity of its

optimization metric to that variable:

$$\text{Care}(x) = \frac{\partial M}{\partial x}.$$

A system cares about  $x$  to the extent that changes in  $x$  move  $M$ . If  $\frac{\partial M}{\partial x} = 0$ , then no optimization pressure—however intense—will preserve  $x$ . The system is not hostile toward  $x$ . It is indifferent. And indifference, at scale, is more destructive than hostility, because it generates no resistance and requires no decision.

The repair infrastructure on which admissibility depends is precisely the class of variables for which  $\frac{\partial M}{\partial x} \approx 0$  in dominant optimization systems. Institutional memory does not move quarterly earnings. Ecological resilience does not move engagement metrics. Tacit expertise does not move market share. Language vitality does not move platform revenue. The result is not that optimization systems attack these goods. It is that they pass through them without registering them, consuming them the way a moving vehicle consumes road surface—not maliciously, not even knowingly, simply as a byproduct of forward motion.

This is why platform scale changes the nature of the problem rather than merely its degree. A small firm with low  $\frac{\partial M}{\partial x}$  for repair goods consumes repair capacity slowly. A system operating at civilizational scale with the same indifference consumes repair capacity at a rate that outpaces the autonomous processes that would otherwise replenish it. The carelessness is the same in kind. The consequence is different in magnitude.

The carelessness principle also explains why appeals to corporate responsibility consistently underperform. Such appeals attempt to increase  $\frac{\partial M}{\partial x}$  through reputational pressure—making the firm care about  $x$  by making  $x$  visible in the metric that governs it. This sometimes works at the margin. It cannot work structurally, because the indifference is not a preference that can be changed by moral argument. It is a property of the optimization landscape. The landscape does not become less indifferent because the people navigating it feel bad about their trajectories.

What would change the landscape is a change in  $M$  itself—which is precisely what the interventions in subsequent sections attempt to describe.

The carelessness principle also provides the unifying formal language for two mechanisms described in earlier sections: chokepoint capitalism and austerity governance. Both can be precisely characterized as institutional mechanisms that force the optimization gradient to become orthogonal to the repair substrate.

**Chokepoints as dimension reduction.** When a platform controls the gateway through which

activity must pass, the system’s optimization metric becomes structurally insensitive to the health of alternative routes. In terms of the care operator, the cross-partial derivatives for all non-gateway pathways are set to zero:

$$\frac{\partial M}{\partial x_{\text{alternative}}} = 0.$$

The alternative route may continue to exist. It generates no signal in  $M$ . No optimization pressure will preserve it. The chokepoint does not destroy alternatives actively; it renders them invisible to the gradient, after which competitive dynamics eliminate them passively.

**Austerity as forced indifference.** Fiscal consolidation mandates that public infrastructure and institutional memory be evaluated strictly on short-term liquid liability. This is equivalent to institutionally enforcing a state where:

$$\frac{\partial M}{\partial x_{\text{repair}}} \rightarrow 0.$$

Maintenance workers, archivists, public health inspectors, and infrastructure engineers generate no return on the balance sheets that govern their employment. Their elimination is not irrational within the governing metric; it is the locally optimal move. Austerity does not choose to destroy repair capacity. It enforces an accounting framework under which repair capacity is invisible, and then allows the gradient to eliminate what the gradient cannot see.

Both chokepoints and austerity are therefore instances of the same structural operation: the deliberate zeroing of  $\frac{\partial M}{\partial x_{\text{repair}}}$  for some class of repair goods. They differ in mechanism—one through platform architecture, one through fiscal policy—but produce the same admissibility consequence: repair capacity is eliminated not through hostile action but through enforced indifference.

## XII. The Admissibility Equation

The preceding sections have described several mechanisms operating simultaneously: compression of the option space through technique; metric reorganization of reality through optimization; structural indifference to repair goods through the carelessness principle; repair deficits through externalization; centralization of distinction governance in platforms; reachability concentration through chokepoint capitalism; deliberate restriction through political governance of admissibility; memory inversion through yield-optimized infrastructure. These are not independent problems. They are related expressions of a single underlying dynamic.

We can represent that dynamic formally. Let  $\mathcal{V}_A(t)$  denote the admissibility volume at time  $t$  the measure of states satisfying  $B(t)x \geq 0$  and recall that the operative quantity governing

civilizational possibility is the admissible reachable set  $\mathcal{R}_A(t) = \mathcal{R}(t) \cap \mathcal{V}_A(t)$ . The dynamics of  $\mathcal{V}_A(t)$  are approximated by:

$$\frac{d\mathcal{V}_A}{dt} = R(t) - C(t) - L(t) - A_P(t),$$

where:

- $R(t)$  is the *repair rate*: the rate at which civilization actively restores the distinctions, institutions, ecological systems, and material infrastructure on which admissibility depends.
- $C(t)$  is *compression-induced distortion*: the rate at which optimization of compressed metrics—and the reorganization of systems around those metrics via  $X' = F(X, \pi(X))$ —diverges trajectories from conditions that preserve admissibility.
- $L(t)$  is *irreversible distinction loss*: the rate at which distinctions, once collapsed, cannot be reconstructed—languages lost, species extinct, institutional knowledge without successors, ecological thresholds crossed, distinction-governing institutions captured by incompatible optimization gradients.
- $A_P(t)$  is *political admissibility restriction*: the rate at which deliberate governance choices—fiscal austerity, institutional restructuring, constitutional constraints—remove regions of future possibility from the reachable state space. Unlike the other terms,  $A_P(t)$  can be negative: governance can also expand  $\mathcal{V}_A$  by rebuilding repair capacity, restoring distributed infrastructure, or opening previously foreclosed pathways.

A critical feature of this system is that the terms interact rather than merely add. Repair capacity is itself degraded by the other mechanisms. We can express this as:

$$R(t) = R_M(t) + R_A(t),$$

where  $R_M(t)$  is metric-driven repair—repair that occurs because it appears in some optimization metric, however imperfect—and  $R_A(t)$  is autonomous repair: the repair activity that persists outside dominant optimization structures because humans engage in it anyway. Volunteers maintain knowledge commons. Archivists preserve institutional memory. Open-source contributors sustain shared infrastructure. Community organizers repair social fabric. Parents transmit cultural distinctions. Teachers maintain epistemic standards. None of these activities are fully explained by optimization pressure; they represent a repair field that survives despite hostile gradients rather than because of favorable ones.

Furthermore,  $R_M(t)$  is itself a function of the other terms:

$$R_M(t) = R_0 - \alpha C(t) - \beta A_P(t),$$

meaning that compression-induced distortion and political restriction both degrade metric-driven repair capacity directly. Substituting:

$$\frac{d\mathcal{V}_A}{dt} = R_0 + R_A(t) - (1 + \alpha) C(t) - L(t) - (1 + \beta) A_P(t).$$

The mechanisms now interact. Compression does not merely add to the contraction directly via  $C(t)$ ; it also subtracts from repair capacity via  $\alpha C(t)$ , amplifying its effect by a factor of  $(1 + \alpha)$ . Austerity governance does not merely restrict admissibility directly via  $A_P(t)$ ; it also destroys repair capacity via  $\beta A_P(t)$ , amplifying its effect by  $(1 + \beta)$ . The metacrisis is not additive. It is compounding.

To fully capture the dynamics described throughout this essay, however, the system must be extended to account for two additional nonlinear feedbacks that convert compounding contraction into potential civilizational phase transition.

First, autonomous repair  $R_A(t)$  depends on the shared epistemic and institutional infrastructure that constitutes the distinction commons—the same infrastructure encoded in  $B(t)$ . As  $\mathcal{V}_A$  contracts toward a critical threshold  $\mathcal{V}_{\text{crit}}$ , this commons degrades and the transmission efficiency of cultural distinctions drops non-linearly. Model this as:

$$R_A(t) = \phi(\mathcal{V}_A(t)),$$

where  $\phi$  is a sigmoidal function:  $\phi(\mathcal{V}_A) \approx R_A^{\text{max}}$  when  $\mathcal{V}_A \gg \mathcal{V}_{\text{crit}}$ , and  $\phi(\mathcal{V}_A) \rightarrow 0$  as  $\mathcal{V}_A \rightarrow \mathcal{V}_{\text{crit}}$ . When the option space contracts too far, the autonomous repair processes that have acted as a persistent brake begin to lose their substrate and collapse toward zero. This is the repair floor collapse: the last buffer against irreversible contraction disappears precisely when it is most needed.

Second, irreversible distinction loss accelerates as the option space narrows. Systems under strain over-optimize and over-specialize, destroying remaining distinctions faster than they would under normal conditions. Model this as:

$$L(t) = L_0 + \frac{k}{\mathcal{V}_A(t)},$$

where  $k > 0$ . As  $\mathcal{V}_A \rightarrow 0$ , the distinction loss rate  $L \rightarrow \infty$ . Incorporating both feedbacks, the full

coupled system is:

$$\frac{d\mathcal{V}_A}{dt} = R_0 + \phi(\mathcal{V}_A) - (1 + \alpha) C(t) - \left( L_0 + \frac{k}{\mathcal{V}_A} \right) - (1 + \beta) A_P(t).$$

This system admits bifurcation behavior. For large  $\mathcal{V}_A$ ,  $\phi(\mathcal{V}_A) \approx R_A^{\max}$  and  $k/\mathcal{V}_A \approx 0$ : the repair brake is active and distinction loss is bounded. For  $\mathcal{V}_A$  near  $\mathcal{V}_{\text{crit}}$ ,  $\phi(\mathcal{V}_A) \rightarrow 0$  and  $k/\mathcal{V}_A$  grows without bound: the repair brake collapses and distinction loss accelerates, driving  $\frac{d\mathcal{V}_A}{dt}$  strongly negative regardless of  $R_0$  or  $A_P(t)$ . The system passes through a critical threshold below which recovery requires not merely policy change but discontinuous external intervention—a structural rupture of the kind the admissibility framework defines as civilizational collapse.

The mathematical significance of this formulation is that it gives precise content to the phrase “civilizational collapse.” It is not a political metaphor. It is the condition  $\mathcal{V}_A(t) \leq \mathcal{V}_{\text{crit}}$  combined with  $\frac{d\mathcal{V}_A}{dt} \ll 0$  at that threshold, producing a self-reinforcing contraction from which the system cannot recover through internal repair alone.

The equation as stated models a system in which admissibility can be restored but not genuinely expanded. Repair returns  $\mathcal{V}_A$  toward prior states; it does not open regions of state space that have never been occupied. Historical expansions, however, often involve genuine novelty: the internet did not restore a lost pathway but opened an entirely new region of  $\mathcal{M}$ . Public libraries, universal education, and open-source software similarly created admissible states that had not previously existed. These are not repair in the strict sense; they are *novelty generation*—the creation of new constraints in  $B(t)$  that add to the admissible set rather than restore degraded ones.

A topological clarification is needed here. Novelty generation operates through two distinct mechanisms with different formal consequences. The first is *boundary reconfiguration*: genuine institutional or technological innovation alters the structure of  $B(t)$  itself, changing its rank or the dimensionality of the coherently occupiable region, effectively opening new directions in  $\mathcal{M}$  that were previously blocked or structurally inaccessible. The invention of the printing press, the development of representative democracy, and the emergence of open-source licensing all reconfigured  $B(t)$  in this sense: they did not merely add states within the existing constraint structure but altered what constraints were possible. The second mechanism is *volume expansion within existing boundaries*: innovation that does not alter  $B(t)$  but adds admissible states within the current structure — a new language being documented, a new cooperative form being registered, a new commons being established. Both are captured by  $N(t)$ , but the first is more fundamental and more difficult to produce, since it requires not just experimentation but the propagation of new constraint structures across the social system.

Formally, this suggests augmenting the system with a novelty generation term  $N(t) \geq 0$ :

$$\frac{d\mathcal{V}_A}{dt} = R_0 + \phi(\mathcal{V}_A) + N(t) - (1 + \alpha)C(t) - \left(L_0 + \frac{k}{\mathcal{V}_A}\right) - (1 + \beta)A_P(t).$$

$N(t)$  represents the rate at which genuinely new admissible pathways are opened—through institutional invention, technological discovery, legal innovation, or the emergence of new forms of social coordination that expand rather than merely restore  $B(t)$ . The distinction between  $R(t)$  and  $N(t)$  matters because their conditions of possibility differ. Repair depends on the persistence of prior distinctions: you can only restore what was, which requires memory of what was. Novelty depends on the capacity to generate new distinctions under conditions that do not yet exist: it requires slack, experimentation space, and the social infrastructure within which peripheral innovation can propagate to the center. Both are suppressed by the current optimization landscape, but through different mechanisms. Repair is suppressed by  $\frac{\partial M}{\partial x_{\text{repair}}} \approx 0$ . Novelty is suppressed by the elimination of the slack and redundancy within which genuine experimentation becomes possible. The cases identified in Sources of Admissibility Expansion—peripheral communities, institutional mutations, open-source ecosystems—are the primary current sites of  $N(t)$ . Their significance lies not only in what they directly contribute to  $\mathcal{V}_A$  but in their role as the generative substrate from which larger expansions have historically emerged.

The metacrisis, stated precisely, is the claim that  $\frac{d\mathcal{V}_A}{dt} < 0$  and that the current trajectory is moving toward rather than away from  $\mathcal{V}_{\text{crit}}$ . It is a framework for asking the question precisely, not a proof that the answer is negative. Its value is that it converts the impressionistic claim that “everything is connected and getting worse” into a falsifiable assertion with identifiable variables, identifiable interaction structures, and identifiable phase transition behavior.

### XIII. What Repair Would Look Like

The standard policy responses to the metacrisis—tighter regulation, better incentives, stronger institutions, more trustworthy media—are not wrong. They are incomplete in a specific way: they optimize within a shrinking admissibility volume rather than addressing the rate of contraction.

The question the admissibility framework asks is different: what interventions increase  $\mathcal{V}_A$  rather than optimizing within it?

**Governing for admissibility expansion.** The  $A_P(t)$  term is the only one in the admissibility equation that can be directly negative—governance can expand  $\mathcal{V}_A$  as well as contract it. Policies that restore slack, rebuild institutional capacity, expand public provision of distinction infrastructure, and reduce the fiscal discipline that suppresses repair capacity are all interventions that

increase  $\mathcal{V}_A$  by reversing deliberate past contractions. The relevant question for any governance intervention is not whether it is fiscally efficient but whether it increases or decreases the range of futures a civilization can still reach. Austerity that destroys repair capacity is not neutral fiscal discipline; it is admissibility contraction by policy. Investment that restores distributed institutional infrastructure is not mere expenditure; it is admissibility expansion by policy.

**Restoring pathway redundancy.** Chokepoint capitalism contracts admissibility by collapsing many reachability routes into a small number of governed passages. The admissibility corrective is not necessarily to eliminate the chokepoints—many provide genuine efficiency—but to maintain alternative routes alongside them. Antitrust policy, interoperability requirements, open protocol mandates, and public infrastructure investment all function, in admissibility terms, as redundancy restoration: they preserve routes that the chokepoint would otherwise eliminate, keeping  $\mathcal{V}_A$  larger than any single gateway’s optimization gradient would permit.

**Internalizing repair costs.** The structural source of the repair deficit is that repair is externalized from optimization metrics. Any intervention that brings repair costs inside the accounting boundary of relevant decision-makers directly increases  $R(t)$ . This is the correct framing for carbon pricing, extended producer responsibility legislation, and mandatory social impact accounting—not as moral requirements but as structural corrections to the divergence  $D = \|M - S\|$ . The goal is not to punish firms. It is to reduce the distance between the metric and the underlying system so that optimization pressure produces fewer inadmissible trajectories.

**Rebuilding distributed distinction infrastructure.** The centralization of distinction governance in a small number of platforms is a direct reduction of  $R(t)$  and increase of  $L(t)$ . Libraries, local journalism, universities, and civil institutions performed distributed distinction repair at a scale and with a redundancy that made the overall system resilient. Their decline is not merely a cultural loss. It is a structural reduction in the civilization’s repair capacity. Rebuilding them—or building successor institutions with equivalent function—increases  $R(t)$  in ways that no amount of platform regulation accomplishes on its own.

**Governing systems rather than metrics.** The Compression Trap closes when the metric becomes the optimization target, and tightens when the system reorganizes around the metric. The corrective is not to find better metrics—this is a trap within a trap, since any sufficiently powerful optimization process will eventually capture any metric—but to govern the underlying systems directly. What matters is not whether GDP is growing but whether the systems GDP was designed to approximate are functioning: whether people have meaningful work, whether ecological

systems remain viable, whether communities have the resources to maintain themselves. Governing toward those substantive conditions, rather than their compressed representations, reduces  $C(t)$ .

**Rebuilding the maintenance class.** Civilization structurally rewards creation over maintenance. Founders are celebrated; repairers are invisible. Investors fund growth; maintenance is a cost center. New products generate press coverage; infrastructure upkeep does not. This is not a cultural accident. It is a structural consequence of the optimization landscape: maintenance does not produce the kinds of measurable output that  $M(x)$  can capture, so maintenance is systematically under-rewarded.

The connection to austerity governance is not incidental. The workers who constitute the maintenance class—teachers, public health workers, inspectors, librarians, transit crews, archivists, infrastructure engineers—are frequently the first expenditures removed during fiscal consolidation. This is the mechanism through which  $A_P(t)$  acts on  $R(t)$ : austerity does not merely remove resources in the abstract; it removes the specific people and institutions whose work is repair. The  $\beta$  coefficient in  $R_M(t) = R_0 - \alpha C(t) - \beta A_P(t)$  is not small. The maintenance class is precisely the population whose labor constitutes  $R_M(t)$ ; their systematic removal during fiscal consolidation is what makes austerity an amplifying mechanism rather than merely an additive one.

Reversing this requires changes to the accounting systems that determine what counts as productive activity—not merely cultural celebration of repair work, but structural recognition that repair work is constitutive of all the other activities that appear in the metrics.

#### XIV. Sources of Admissibility Expansion

The analysis to this point explains contraction. A critic will rightly observe that an account of how traps are constructed does not explain how they are escaped—and that if optimization gradients are as powerful as described, the interventions listed in the previous section face a fundamental problem: they must be implemented by agents who are themselves embedded in the optimization landscapes they are trying to change.

This objection is serious. But it has a historical answer. Admissibility has been expanded before, not by escaping optimization entirely but through specific mechanisms that operate alongside and sometimes against dominant gradients. Identifying these mechanisms transforms the analysis from diagnosis to theory.

**Crisis-generated reorganization.** Systems under sufficient stress sometimes undergo rapid reorganization that opens new option spaces rather than closing them. The New Deal, postwar

reconstruction in Western Europe, the development of national health systems, the environmental legislation of the 1970s—all occurred in the aftermath of crises that delegitimized the optimization landscapes that had produced them. Admissibility expansion does not require that people escape the gradient; it requires that the gradient itself be discredited sufficiently to permit institutional redesign. The crisis does not automatically produce good outcomes—it also produces fascism, authoritarianism, and war—but it creates the conditions under which  $A_P(t)$  can be negative at a scale sufficient to matter.

**Institutional mutation.** Existing institutions occasionally develop internal repair functions that exceed what their optimization metrics would predict. Regulatory agencies sometimes enforce against the industries that nominally capture them. Scientific institutions sometimes produce findings that undermine the commercial interests funding them. Legal systems sometimes generate precedents that constrain the powerful. These outcomes are not random; they tend to occur where institutional cultures have successfully defended some degree of autonomy from the dominant optimization gradient. They are examples of  $R_A(t)$ —autonomous repair that operates despite rather than because of the surrounding incentive structure. The question of how to protect and extend such autonomy is a specific institutional design problem, not an abstract moral one.

**Peripheral experimentation.** The dominant optimization landscape does not cover all social space simultaneously. At the periphery—in open-source communities, mutual aid networks, indigenous land management systems, cooperative enterprises, academic subfields insulated from commercial pressure, small cities experimenting with governance, community archives, amateur preservation societies—alternative approaches to repair, maintenance, and distinction governance persist and sometimes flourish. These peripheries are important not only for their direct repair contributions but as existence proofs: they demonstrate that  $R_A(t) > 0$ , that repair continues outside the dominant gradient, and that some institutional forms remain available as models for larger-scale reconstruction.

**Repair coalitions.** Admissibility expansion historically tends to require coalitions whose members are motivated by different and sometimes incompatible reasons. The coalition that produces environmental regulation may include scientists motivated by evidence, communities motivated by health harms, activists motivated by values, and businesses motivated by competitive advantage from cleaner technology. The coalition that produces labor protections may include unions motivated by wages, reformers motivated by dignity, and managers motivated by reduced turnover. These coalitions do not require that all members share the same metric. They require only that enough members share a direction—that  $\mathcal{V}_A$  should increase rather than decrease—and that their

combined capacity to act exceeds the capacity of those benefiting from contraction to prevent it.

The historical record does not support optimism about any particular outcome. Empires contract. Institutions fail. Repair coalitions dissolve. But it supports something more useful than optimism: the observation that admissibility expansion is a recurring possibility, not a structural impossibility. The question the admissibility framework poses is not whether expansion is guaranteed, but what conditions make it more or less likely—and what can be done, within existing constraints, to move toward those conditions.

## XV. Toward Measurement

A framework that cannot be falsified in principle is philosophy. A framework that cannot be falsified in practice is premature theory. The admissibility framework currently occupies the second category: the formal structure is in place, but the variables are not yet operationalized. Closing that gap is the work of subsequent research, but the present essay can at least specify what measurement would require.

The central quantity is the admissible reachable set  $\mathcal{R}_A(t) = \mathcal{R}(t) \cap \mathcal{V}_A(t)$ . Its contraction proceeds through two channels: degradation of  $B(t)$ , which shrinks  $\mathcal{V}_A(t)$ , and concentration of pathways, which shrinks  $\mathcal{R}(t)$ . Proxies for each must be tracked separately.

The boundary operator  $B(t)$  encodes the distinction infrastructure on which admissibility depends. Its degradation is what drives  $\mathcal{V}_A$  contraction. Proxies for  $B(t)$  degradation are not hard to identify even if they are imperfect: rates of institutional closure and consolidation; linguistic diversity indices tracking language vitality; epistemic trust measures from longitudinal surveys; ecological resilience indicators such as species diversity, soil health, and aquifer levels; archival completeness metrics; journalistic capacity measures (local newsroom counts, investigative output rates); and rates of knowledge domain extinction as measured by citation network fragmentation.

The repair rate  $R(t)$  is approximable through: infrastructure maintenance expenditure as a share of GDP; institutional longevity distributions; volunteer activity indices; open-source maintenance rates and contributor counts; civic participation measures; and the relative compensation of maintenance-class occupations versus creation-class occupations.

Compression distortion  $C(t)$  is approximable through: the spread between financialization indices and underlying productive capacity measures; the gap between engagement metrics and self-reported communicative satisfaction; the divergence between educational attainment metrics and skill assessment outcomes; and corporate expenditure on representational repair (public relations, lobbying, narrative management) versus substrate repair (R&D, infrastructure, workforce

development).

Irreversible distinction loss  $L(t)$  is approximable through: language extinction rates; species extinction rates; institutional memory loss as measured by policy rediscovery cycles (how frequently policy failures are repeated); and the rate at which tacit expertise is lost to retirement without successor training.

Political admissibility restriction  $A_P(t)$  is approximable through: austerity indices tracking public investment ratios; institutional closure rates during fiscal consolidation; labor protection indices; and the fiscal space available for repair-intensive public goods.

The novelty generation term  $N(t)$  is the most difficult to proxy, because novelty by definition opens previously nonexistent regions of  $\mathcal{M}$ . Approximations include: rates of genuinely new institutional form adoption (cooperative enterprises, new legal structures, new forms of commons governance); open-source project creation rates in domains not previously covered; rate of new language vitalization efforts; and the density of peripheral experimentation in governance, technology, and social organization. Declining  $N(t)$  proxies alongside declining  $R(t)$  proxies would provide evidence that the system is losing not only its capacity to maintain existing pathways but its capacity to generate new ones.

None of these proxies directly measure  $\mathcal{V}_A$ . But tracking them simultaneously, looking for correlated degradation across domains, would provide evidence bearing on the central claim that  $\frac{d\mathcal{V}_A}{dt} < 0$  and that the mechanisms of contraction are coupled rather than independent. The framework makes a specific prediction: the proxies should correlate with each other more strongly over time, because the coupling terms  $\alpha$  and  $\beta$  mean that contraction in one domain accelerates contraction in others. If the proxies are largely uncorrelated, the coupling hypothesis is wrong.

The threshold  $\mathcal{V}_{\text{crit}}$  deserves particular care. It is the point below which autonomous repair collapses and distinction loss becomes self-accelerating. Identifying it empirically would require evidence of phase transition behavior: sudden non-linear acceleration in the rate of distinction loss coinciding with collapse in autonomous repair activity. Historical cases—the late Western Roman Empire, the late Qing administrative system, the collapse of Soviet institutional infrastructure—provide potential data. They also provide a warning: the threshold is typically identified only in retrospect. This is precisely why the framework’s value lies not in predicting collapse but in identifying the conditions under which the threshold becomes more or less distant.

One political implication of  $\mathcal{V}_{\text{crit}}$  deserves explicit acknowledgment, and it involves a secondary feedback that sits squarely within the Baudrillardian layer of the framework. The concept could be weaponized: a claim that “we have crossed the threshold” could be used to justify suspension

of normal democratic process in the name of emergency admissibility restoration. This danger is not hypothetical; it is the standard structure of authoritarian emergency claims. But there is a more subtle version of the same problem. The *representation* of approaching  $\mathcal{V}_{\text{crit}}$ —a forecast, a declaration of crisis, an expert determination that the threshold is imminent—can itself function as a positive  $A_P(t)$  term, artificially contracting the option space under the guise of preservation. This is the emergency paradox: the prediction of admissibility collapse becomes a mechanism of admissibility collapse, as institutions respond to the representation of crisis by closing down the political processes through which expansion would otherwise become possible. The representation of the threshold becomes constitutive of the threshold itself.  $X' = F(X, \pi(X))$  applies not only to economic metrics but equally to political crisis narratives. The admissibility framework is therefore subject to its own central dynamic: its representations can be captured, and a captured admissibility framework would be used to justify precisely the option-space closures it was designed to identify and resist. The appropriate response is not to abandon the threshold concept but to insist that the identification of  $\mathcal{V}_{\text{crit}}$  requires democratic deliberation rather than expert determination, and that a framework designed to expand the space of legitimate futures cannot coherently be invoked to foreclose it.

## XVI. Conclusion: The Prior Crisis

A civilization does not collapse when it exhausts its resources. It collapses when it exhausts its capacity to reach states in which its resources can still be used.

The admissibility crisis is prior to the resource crisis, prior to the attention crisis, prior to the trust crisis, prior to the ecological crisis. It is the condition under which all the others become unsolvable—not because solutions do not exist, but because the structural capacity to implement and sustain them has been progressively consumed.

This is what distinguishes the admissibility framework from standard metacrisis analysis. Standard analysis identifies a list of failures and argues that they are connected. The admissibility framework identifies a single variable—the rate of change of the option space—and shows that the standard failures are expressions of its contraction. The distinction matters because it changes the target of intervention. You cannot fix a repair deficit by addressing any particular thing being repaired. You have to address the conditions under which repair is possible.

Ellul warned that societies organized around technique progressively subordinate all competing values to efficiency, generating self-expanding optimization landscapes that consume option space to produce local performance. Baudrillard warned that societies organized around representations progressively lose contact with what those representations once described—and eventually find

that the representations have become constitutive of the reality they claimed merely to measure. Doctorow observed that the most strategically significant positions in modern economies are not productive but positional—governing the pathways through which all other activity must pass, concentrating reachability in the hands of gateway operators. Mattei demonstrated that admissibility contraction is not always emergent: political institutions have historically and deliberately narrowed the range of futures that remain reachable, presenting the narrowing as technical necessity while it functioned as political project. Together these four accounts describe a single converging dynamic: technique compresses what can be measured; simulation reorganizes reality around the measurement; chokepoints concentrate the routes through which reorganized reality can be navigated; and governance removes the destinations that redistribution of navigational capacity would otherwise make available.

One question the framework raises but does not answer is the relationship between economic growth and admissibility. The analysis might seem to imply that growth is structurally incompatible with admissibility preservation—that the optimization gradient is inherently contractionary and that expansion of productive capacity therefore inherently contracts the option space. This is not the framework’s position. The problem is not growth per se but the specific optimization gradient under which growth currently occurs: one in which  $\frac{\partial M}{\partial x_{\text{repair}}} \approx 0$  for the goods that maintain  $B(t)$ . Growth directed toward repair—investment in maintenance infrastructure, ecological restoration, institutional capacity, distributed knowledge commons—would increase  $R(t)$  and expand  $\mathcal{V}_A$ . The framework is therefore compatible with growth as such while being fundamentally incompatible with growth under an optimization gradient that is structurally indifferent to the conditions that make growth sustainable. The distinction matters: if growth itself is the problem, the required intervention is more radical than anything this essay proposes. If the optimization gradient is the problem, the required intervention is a change in what counts as productive—which is precisely what the essay argues.

The metacrisis is not the accumulation of failures. It is the progressive contraction of the admissible future under conditions where the maps have become more valuable than the territories they were meant to navigate.

The reversal requires something specific: not a better map, not a cleaner optimization target, not a more efficient repair process. It requires a change in what counts as productive. It requires that maintenance, repair, distinction preservation, and institutional memory—the activities that produce admissibility rather than output—be brought inside the accounting boundary where they can appear in the metrics that drive behavior.

Until they do, the gradient points the same direction it has been pointing. The volume contracts.

The options narrow. Not because anyone chose this. Because no one is measuring what it costs.

## References

- [1] Arrow, Kenneth J. *Social Choice and Individual Values*. 2nd ed. New Haven: Yale University Press, 1963.
- [2] Baudrillard, Jean. *Simulacra and Simulation*. Translated by Sheila Faria Glaser. Ann Arbor: University of Michigan Press, 1994.
- [3] Berkes, Fikret, Johan Colding, and Carl Folke. *Navigating Social–Ecological Systems: Building Resilience for Complexity and Change*. Cambridge: Cambridge University Press, 2008.
- [4] Doctorow, Cory. *Chokepoint Capitalism*. New York: Beacon Press, 2023.
- [5] Ellul, Jacques. *The Technological Society*. Translated by John Wilkinson. New York: Vintage Books, 1964.
- [6] Foucault, Michel. *Discipline and Punish: The Birth of the Prison*. New York: Pantheon Books, 1977.
- [7] Graeber, David. *Bullshit Jobs: A Theory*. New York: Simon & Schuster, 2018.
- [8] Hirschman, Albert O. *Exit, Voice, and Loyalty*. Cambridge, MA: Harvard University Press, 1970.
- [9] Illich, Ivan. *Tools for Conviviality*. New York: Harper & Row, 1973.
- [10] Jacobs, Jane. *The Death and Life of Great American Cities*. New York: Random House, 1961.
- [11] Kauffman, Stuart A. *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press, 1993.
- [12] Latour, Bruno. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford: Oxford University Press, 2005.
- [13] Mattei, Clara E. *The Capital Order: How Economists Invented Austerity and Paved the Way to Fascism*. Chicago: University of Chicago Press, 2022.
- [14] Meadows, Donella H. *Thinking in Systems: A Primer*. White River Junction, VT: Chelsea Green Publishing, 2008.
- [15] North, Douglass C. *Institutions, Institutional Change and Economic Performance*. Cambridge: Cambridge University Press, 1990.
- [16] Ostrom, Elinor. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press, 1990.

- [17] Perrow, Charles. *Normal Accidents: Living with High-Risk Technologies*. New York: Basic Books, 1984.
- [18] Polanyi, Karl. *The Great Transformation*. Boston: Beacon Press, 2001.
- [19] Putnam, Robert D. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon & Schuster, 2000.
- [20] Scott, James C. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven: Yale University Press, 1998.
- [21] Shannon, Claude E. "A Mathematical Theory of Communication." *Bell System Technical Journal* 27, no. 3 (1948): 379–423.
- [22] Simon, Herbert A. "The Architecture of Complexity." *Proceedings of the American Philosophical Society* 106, no. 6 (1962): 467–482.
- [23] Taleb, Nassim Nicholas. *Antifragile: Things That Gain from Disorder*. New York: Random House, 2012.
- [24] Tainter, Joseph A. *The Collapse of Complex Societies*. Cambridge: Cambridge University Press, 1988.
- [25] Winner, Langdon. *The Whale and the Reactor: A Search for Limits in an Age of High Technology*. Chicago: University of Chicago Press, 1986.