

Structure Through Change

Entropy, Distinction, Repair, and the Making of Persistent Worlds

Flyxion
Independent Researcher

September 2026

Abstract

Nearly everything that gets called a thing is, on inspection, a process that has not yet stopped. A river keeps no water. An organism keeps almost none of its molecules for long. A star burns through its own composition as a condition of remaining a star. A conversation, a program, an institution, a body of research: each survives by replacing its own material at every scale while some more abstract organization carries over. This essay treats that carrying-over as the object of study in its own right, rather than as an incidental feature of otherwise separate domains. It develops a single trajectory, entropy to distinction to structure to persistence to repair to continuation, and shows that this trajectory, not any one physical or computational model, is what unifies a body of work spanning field theory, mereology, historical reconstruction, cognitive architecture, semantic theory, programming-language design, biological aging, agentic delegation, and institutional design. The claim is not that these domains reduce to one another. It is that they answer regime-specific versions of one question: what has to remain organized for something to go on being itself while nothing about it stays the same twice. The essay develops the trajectory in full, traces its application across a range of regimes, physical, mereological, historical, cognitive, semantic, computational, biological, and institutional, without claiming these categories are mutually exclusive or that the count among them is itself a finding rather than an organizational choice, sketches the mathematical vocabulary the trajectory borrows and extends, and closes by isolating one consequence that reorders the usual relation between persistence and truth.

Contents

1	The problem, stated plainly	2
2	Entropy: the background condition	3
3	Distinction: before there is a candidate for persistence	4
4	Structure as constraint, and its regimes	6
5	A physical substrate for localized persistence	7
6	Reconstruction after loss	9
7	Repair: persistence surviving disruption	10

8 Continuation: history as a constraint on the future	12
9 Cognitive and semantic continuation	14
10 Computation as transformation ontology	16
11 Biology and longevity	17
12 Delegation, recovery margin, and contestability	18
13 The mathematical vocabulary of invariance	20
14 The trajectory across regimes: a synthesis	21
15 Persistence as a condition on truth	21
16 Open problems	22
17 Conclusion	23

1 The problem, stated plainly

The ordinary way of thinking about identity starts from an object and treats change as something that happens to it afterward. A chair is a chair, and then it gets scratched, and it is still a chair. That framing works well enough for chairs because chairs change slowly and the parts that change are not the parts anyone cares about. It works far worse for a river, an organism, a habit, a company, a conversation, or a star, because in each of these cases the changing material is not incidental to the thing. It is close to the whole of what the thing is made of, moment to moment. A river is made of whatever water happens to be passing through the riverbed right now. An organism is made of a population of molecules with individually short residence times, most of them replaced within weeks or months. A conversation is made of turns that vanish into a transcript the moment they are spoken. A company is made of whichever employees currently hold its roles, none of whom need have been present at its founding. None of these things has a stable material core waiting underneath the flux, and the felt temptation to look for one, to imagine that if only the scratched surface of the chair were peeled back far enough there would be a permanent something at the center, is exactly the intuition this essay sets out to dismantle.

That observation forces a reversal. Instead of asking what stays the same while an object is disturbed by change, the more tractable question is what kind of organization is capable of surviving as an ongoing achievement of change, rather than as a possession held in spite of it. This essay calls that achievement *structure*, and its temporally extended survival *persistence*, and takes the position that persistence is not the default condition of anything. It is manufactured, continuously, by systems that happen to be organized in a way that makes manufacture possible. When the manufacturing fails locally, the system either falls apart, turns into something else, or is repaired back into an admissible version of what it was. When the manufacturing succeeds over a long run, the system's present state comes to depend on its own history in a specific, formalizable way: some future states remain reachable and some do not, and that partition is the content of what this essay calls *continuation*.

The argument runs through six ideas in sequence: entropy, distinction, structure, persistence, repair, and continuation. None of the six is proprietary to one discipline. Entropy belongs to ther-

modynamics, distinction to mereology and to logic, persistence to metaphysics and to biology, repair to engineering and to medicine, continuation to narrative theory and, more formally, to the theory of constrained dynamical systems. What is proprietary here is the claim that they form one line of development rather than six separate topics, and that a wide range of existing frameworks, a field-theoretic account of localized organization, a scale-indexed mereology, a theory of reconstruction after loss, an account of repair warrants, and a formal theory of continuation operators, are successive elaborations of that one line rather than an assortment of unrelated projects loosely joined by shared vocabulary.

Three clarifications forestall likely misreadings before the argument begins. First, none of what follows denies that objects exist, or that some things really are more stable than others. A boulder genuinely persists differently than a candle flame, and nothing in the framework below erases that difference. What the framework denies is that the boulder's stability is metaphysically prior to the flame's, as though the flame were a defective or merely apparent case of persistence and the boulder the pure case. Both are achievements of organization under transformation; they differ in the rate and visibility of the transformation involved, and in almost nothing else that matters to the present argument. Second, treating persistence as manufactured rather than default is not equivalent to treating it as illusory. A river is not less real for depending on continuous inflow; the dependency is part of what a river is, not a debunking of the concept. Third, the essay does not claim that every domain surveyed here, physics, biology, cognition, institutions, reduces to the others. The recurring trajectory is a shared question, answered differently and with different formal tools in each regime, not a hidden identity between the regimes themselves.

2 Entropy: the background condition

Structure cannot be discussed sensibly without a background against which its survival is non-trivial. That background is supplied by processes that tend to erase organization: mixing, decay, dissipation, forgetting. Call the generic tendency entropy, without committing yet to any particular physical formalism, since the argument needs only the qualitative fact that distinguishable arrangements tend, absent intervention, to become less distinguishable from their neighbors over time. The point of naming this tendency first is not to explain persistence by appeal to entropy, since entropy on its own explains only the opposite. The point is to establish that persistence, wherever it occurs, occurs against resistance. A structure that persists is not a structure to which nothing is trying to happen. It is a structure that is winning, locally and temporarily, an argument it is otherwise disposed to lose, and every one of the frameworks discussed later inherits this feature: repair is only intelligible against a background in which damage is the default outcome of time passing, and continuation is only informative because most conceivable futures are ones in which the relevant organization has simply dissolved.

This has an immediate and somewhat counterintuitive consequence: persistence is frequently achieved through continued activity rather than through inertness. A main-sequence star does not persist by being left alone. It maintains its particular quasi-equilibrium through continuous energy generation and transport, converting hydrogen to helium against its own gravitational collapse, and the cessation of core hydrogen fusion does not simply erase the star; it forces a transition into another evolutionary regime, a red giant, a white dwarf, a neutron star, depending on mass and prior history, some of which remain structurally organized for a very long time without any ongoing fusion at all. What the main-sequence case illustrates is narrower than "no activity, no persistence" as a universal law; it illustrates that a given regime of organization, this specific quasi-equilibrium rather than persistence in general, is maintained by continued activity, and that when

the activity stops, the system does not freeze in place but moves to whatever regime its remaining physics permits. An organism does not persist by resisting metabolism. It persists because of metabolism, and it dies the moment metabolism stops, at which point its material composition changes almost not at all in the first instants and yet everything about its status as a persisting organism has already been lost. A well-run institution does not persist by preventing turnover of its staff. It persists partly through turnover, provided the turnover is structured so that institutional memory outlives any individual member, and an institution that succeeded in freezing its membership permanently would not thereby become more persistent; it would become brittle, since it would lose the capacity to replace failing components that every long-lived organization eventually needs.

In each case, flux is not the enemy of the persisting structure. Flux is the medium in which the structure's persistence is enacted. This is the first departure from naive intuition that the rest of the essay depends on: identity is not what survives change, in the sense of standing apart from it. Identity, where it exists at all in these systems, is a pattern maintained *by* change, and would not exist without it. The relevant contrast is not between change and no change, but between two kinds of change: change that disperses organization and change that regenerates it. A useful, if informal, way of marking the distinction is to note that dissipative structures, to borrow a phrase from far-from-equilibrium thermodynamics, require a continuous throughput of energy or material precisely because they are organized; the organization is what has to be paid for, continuously, rather than something that, once achieved, is free to persist on its own. A flame, a whirlpool, a living cell, and by extension the more abstract structures discussed later in this essay, share this feature: cut off the throughput and the structure does not merely weaken, it disappears on a timescale set by its own internal relaxation dynamics, which is typically far shorter than the timescale on which it had appeared stable while the throughput was maintained.

Entropy alone, however, cannot supply the vocabulary needed to say which arrangements count as the same structure across a sequence of transformations. A thermodynamic account can say that a gradient is dissipating, that a distribution is approaching uniformity, that free energy is decreasing. It has nothing to say about why a particular whirlpool in a stream counts as one persisting object rather than a series of unrelated fluid configurations that happen to look similar from one moment to the next, nor why a given organism counts as the same organism from birth to death despite total molecular turnover, nor why an institution counts as the same institution across a change of every officer and every building it occupies. Entropy sets the stakes of the inquiry. It does not, by itself, supply the unit of analysis. That question belongs to the next stage.

3 Distinction: before there is a candidate for persistence

Before it makes sense to ask whether something persists, it has to be settled what would count as the same thing continuing versus a different thing replacing it. That is a question about distinction, not about time. A configuration must first be distinguishable from its surroundings, along some criterion, before its temporal career becomes a meaningful object of inquiry at all; without a prior answer to what counts as the thing, the question of whether the thing persisted cannot even be posed coherently, since there would be no fixed referent for the question to be about.

The criterion need not be spatial. A whirlpool is spatially distinguishable from the rest of the stream, but a species is distinguished by reproductive and genealogical relations rather than by any boundary in physical space, and a semantic category is distinguished by a pattern of inferential relations rather than by any boundary at all in the ordinary sense. A melody is distinguished from the surrounding silence and from other melodies by a pattern of pitch and interval relations

that survives transposition into a different key, a different instrument, and a different tempo, none of which preserve a single acoustic property in common between the original performance and the transposed one. What all of these share is a relation, call it $D(X, E)$, between a candidate structure X and an environment E , such that some transformations preserve the relation and others erase it. A whirlpool survives being nudged by a passing branch and does not survive being stirred vigorously across its whole extent. A species survives ordinary genetic drift and does not survive being defined out of existence by a taxonomic revision that redraws its boundary around a different set of populations. A melody survives transposition and does not survive having its interval pattern scrambled, even if every individual note remains drawn from the same scale. In each case, distinction is not a brute given. It is itself a claim about which transformations are admissible without destroying the thing distinguished, which quietly imports the vocabulary of persistence into the definition of distinction. The two notions are, in this sense, co-defined rather than sequential; entropy sets the stakes, but distinction and persistence arrive together, since to distinguish a structure at all is already to commit to some standard of what would count as its continuing.

It is worth cataloguing, briefly, the range of boundary types that satisfy this general relation, since the variety is easy to underestimate if one starts from the paradigm case of a spatial boundary. A physical boundary, a membrane, a surface, a phase interface, separates a system from its environment by restricting the transport of matter or energy across it, and the degree of restriction, rather than its mere presence, determines how sharply the system counts as distinct. A dynamical boundary, of the kind that distinguishes a vortex from the surrounding fluid, is not a material barrier at all but a region of coherent, correlated motion embedded in a larger, less correlated flow; the vortex has no membrane, only a statistical signature that persists even as the specific fluid parcels inside it are continuously exchanged with the parcels outside. An informational boundary, of the kind at stake in the notion of a Markov blanket, separates a system from its environment by restricting which variables can influence which: internal states are shielded from direct external influence and instead interact with the environment only through a smaller set of sensory and active states at the boundary itself, and this kind of boundary can be present with no accompanying physical barrier whatsoever. A functional or interface boundary, of the kind relevant to engineered systems, separates a component from the rest of a system by specifying a contract, an interface, that other components must respect, so that the internal implementation behind the interface can be replaced wholesale without the rest of the system detecting any difference, provided the interface's guarantees are honored. An institutional or jurisdictional boundary separates a domain of authority or applicability by convention and enforcement rather than by anything physical, dynamical, or informational in the narrower senses just given, and yet it satisfies the same general relation: some transformations, an election, a change of location, a change of personnel, are compatible with the institution's continuity, and others, a dissolution, a hostile takeover that discards the institution's charter, are not.

This matters because it rules out an appealing but mistaken shortcut: the idea that one could first carve the world into objects using some neutral, history-independent criterion, and only afterward ask which of those objects persist. Boundaries drawn without reference to dynamics tend to be boundaries drawn around momentary appearances, and momentary appearances are exactly what fails to track anything interesting about survival through transformation. A boundary earns its keep only if it tracks a regime in which some family of transformations is compatible with continued identity and some is not. That regime-relativity is the seed of the next section, and it also explains why disputes about where, precisely, to draw a boundary, the exact edge of a species, the exact moment an institution ceased to be its former self, are frequently disputes that cannot be settled by finer-grained observation of the present moment alone, because the boundary in question was never a fact about a single instant to begin with.

4 Structure as constraint, and its regimes

Once distinction is understood as tied to admissible transformation rather than to static shape, structure can be defined economically: a structure is a restriction on the space of transformations a configuration can undergo while remaining the same distinguishable thing. Write the restriction as an admissible set $\mathcal{A}(X) \subseteq \mathcal{T}$, where $\mathcal{T}$ is the full space of transformations available in principle. A rigid object has a narrow $\mathcal{A}(X)$: almost nothing can happen to it without it ceasing to be recognizably itself. A river has a much broader $\mathcal{A}(X)$: essentially all of its water can be replaced, its course can shift gradually, its flow rate can vary by orders of magnitude across a year, and it remains the same river under every ordinary description. The size and shape of $\mathcal{A}(X)$ is therefore not a nuisance parameter to be estimated once the real definition of X is otherwise in hand. It is close to being the content of what kind of thing X is, since two systems with identical present material composition but differently shaped admissible sets are, for the purposes of this framework, different kinds of thing.

This licenses a scale-indexed treatment of wholes rather than a single universal criterion of parthood. At the molecular scale, a cell's boundary is a lipid membrane subject to a fairly narrow admissible family of chemical transformations, most of which the membrane actively excludes rather than merely happening to lack. At the organismal scale, admissibility is governed by physiological regulation, and almost the entire lower-level material can turn over while the organism persists, since the invariants that matter at this scale, gross morphology, functional integration among organs, continuity of the regulatory processes that maintain internal conditions within survivable bounds, are compatible with essentially unrestricted turnover at the molecular scale beneath them. At the institutional scale, admissibility is governed by charters, norms, and legal continuity, and every member can be replaced while the institution persists as the same legal and social entity, since the relevant invariant is a pattern of roles, procedures, and recognized authority rather than any fact about which individuals currently occupy those roles. There is no privileged scale at which the real criterion of identity lives, and no obligation to reduce the institutional case to the molecular one. What is required instead is a family of admissibility relations indexed by regime, $\{\mathcal{A}_s\}_{s \in S}$, together with an account of how regimes stack: an institution's admissible transformations are not simply the union of its members' admissible transformations, because institutional continuity survives transformations, wholesale replacement of personnel, for instance, that would not count as identity-preserving at the level of any individual member, and conversely an institution can be destroyed by a transformation, a change of governing charter, that need not touch any individual member's continuity at all.

This scale-indexation avoids two failure modes that recur in less careful treatments of identity-through-change. The first is naive reductionism, the assumption that higher-level persistence is nothing but the sum of lower-level persistences, which cannot explain why an organism outlives essentially all of its constituent molecules many times over, nor why a piece of software can be recompiled for an entirely different processor architecture, with none of its original machine-level instructions surviving in any recognizable form, and still count as the same program. The second is ontological fragmentation, the assumption that each scale is an independent reality with no relation to the others, which cannot explain why disrupting the molecular scale severely enough reliably disrupts the organismal scale as well, nor why corrupting enough of a program's source code reliably breaks its behavior at the level users actually interact with. The right relation is that each regime supplies its own admissibility criterion, but the regimes are coupled: damage propagates upward probabilistically rather than deterministically, and repair, discussed below, is frequently a matter of intervening at one scale to restore admissibility at another, which is exactly what happens when a surgeon operates at the tissue scale to restore function at the organismal scale, or

when an engineer patches a single subroutine to restore correct behavior at the scale of the whole application.

Two further refinements belong here because later sections depend on them. The first concerns closure. A natural whole achieves what can be called natural closure: its boundary and its persistence are maintained by processes intrinsic to the system itself, as in a cell's maintenance of its own membrane potential or an organism's maintenance of its own body temperature. An engineered whole, by contrast, typically achieves engineered closure: its boundary is maintained by a specification external to the physical process realizing it, an interface contract, a protocol, a legal definition, so that the same functional identity can be realized by radically different physical substrates without any loss of continuity, in a way that would be unintelligible for a naturally closed system, where a wholesale substrate change of comparable magnitude would typically be indistinguishable from destruction. This distinction, natural versus engineered closure, is what makes it coherent to say that a piece of software, a corporation, or a governmental office can survive complete replacement of its underlying material in a way that would kill an organism subjected to an analogous replacement rate, since the mechanism enforcing the boundary in the engineered case does not itself depend on any particular physical realization persisting.

The second refinement concerns the observer. Not every partition an observer finds convenient corresponds to a dynamically maintained whole. Represent observation generically as a mapping $\mathcal{O} : \mathcal{P} \rightarrow \mathcal{R}$ from physical configurations into representational structures, a coarse-graining that discards most of the microscopic detail of $\mathcal{P}$ in favor of a smaller set of variables in $\mathcal{R}$ that the observer treats as meaningful. The coarse-graining an observer chooses is an *epistemic* partition; it need not coincide with an *ontic* partition, a dynamically stable regime that would maintain its own boundary regardless of whether anyone were watching. The two coincide, and produce genuinely informative descriptions, exactly when the observer's chosen variables track an attractor of the underlying dynamics rather than an arbitrary slice through it; a thermodynamic description of a gas in terms of pressure, volume, and temperature is a paradigm case of alignment, since these coarse variables track quantities that are themselves dynamically stabilized by the underlying molecular chaos, while a description of the same gas in terms of, say, the total momentum of every particle whose label happens to be a prime number would be an epistemic partition with no corresponding ontic counterpart, since nothing in the gas's dynamics treats that subset as special. Much of what separates a productive scientific ontology from an unproductive one, on this account, is exactly this question of alignment, and it recurs, in a more pointed form, in the discussion of semantic structure below.

5 A physical substrate for localized persistence

The mereological picture above is deliberately substrate-neutral, but it is worth showing that it is not merely a redescription with no physical content, by sketching how localized physical structure can be treated as a special case. Consider a physical system described by a coupled scalar, vector, and entropic field configuration, $(\Phi, \mathbf{v}, S)$, evolving under some dynamics that regulates gradients rather than allowing them to grow or decay without bound. On this picture, an object is not a primitive constituent of the field theory. It is a dynamically stabilized configuration, a region in which the joint evolution of $(\Phi, \mathbf{v}, S)$ satisfies some closure condition, $\mathcal{F}(\Phi, \mathbf{v}, S) \approx 0$, over a timescale long compared to the local relaxation time of the field.

This has the immediate advantage of making precise a distinction that ordinary talk of objects leaves implicit: the distinction between a configuration that is merely momentarily present and one that is dynamically favored. A momentary fluctuation in $(\Phi, \mathbf{v}, S)$ can satisfy the closure condition

by accident, for an instant, without being stabilized by anything in the surrounding dynamics; such a fluctuation disperses on the relaxation timescale and never rises to the level of an object in any useful sense. A configuration that persists over many relaxation times, by contrast, is doing so because the surrounding dynamics actively maintains it, pushing nearby perturbed configurations back toward the closure condition rather than letting them wander away from it. This is the field-theoretic analogue of the admissible set $\mathcal{A}(X)$ introduced above: the closure condition, together with the dynamics that restores it after small perturbations, defines exactly which transformations a given configuration can undergo while remaining the same persistent structure.

This reframing has a further, specific technical virtue: it explains why persistent structure can arise from processes that are locally unstable rather than locally stable, which is otherwise a puzzling feature of many pattern-forming systems. Consider a homogeneous field configuration that is unstable to small perturbations, in the sense that some Fourier modes grow rather than decay, $\omega(k) > 0$ for k in some band, where $\omega(k)$ gives the growth rate of a perturbation with wavenumber k . Unbounded growth of that kind produces nothing durable; it either saturates catastrophically, running away until some other physical limit intervenes, or disperses once nonlinear effects outrun the process that was amplifying the perturbation in the first place. But growth that is bounded to a finite band of wavenumbers, of the schematic form

$$\omega(k) \sim \alpha k^2 - \kappa k^4,$$

initially selects a preferred length scale, set approximately by the wavenumber at which $\omega(k)$ is maximized. This functional form is familiar from the study of spinodal decomposition in phase-separating mixtures, where a homogeneous mixture quenched into an unstable regime spontaneously separates into a pattern of domains at a characteristic early size, the instability regulated at short wavelengths by a fourth-order term that penalizes sharp gradients, and the same qualitative behavior recurs across a wide range of pattern-forming physical systems that otherwise share little else in common.

It is important not to overstate what this selection accomplishes on its own. In ordinary Cahn–Hilliard dynamics, the linear instability sets an early characteristic scale, but it does not by itself lock that scale into a permanently stable pattern: domains generically continue to coarsen after the linear regime saturates, driven by the same free-energy reduction that produced the initial separation, so that the pattern present shortly after the instability first develops is not, in general, the pattern present indefinitely afterward. Persistent, finite-scale organization therefore requires something beyond bounded instability alone. It requires either restricting the claim of persistence to a specified observational interval, over which coarsening has not yet proceeded far enough to matter, or identifying an additional arresting mechanism, competing long-range interactions, external forcing, pinning by disorder or by a substrate, reaction terms that stabilize a finite domain size, or coupling to a second field that penalizes further coarsening, that halts the coarsening process at a finite scale rather than letting it run to completion.

This qualification sharpens rather than weakens the framework’s own logic. Bounded instability explains how differentiated structure can emerge from a homogeneous background. It does not, by itself, explain why that structure should go on existing once it has emerged, and treating emergence and long-term persistence as the same achievement would blur exactly the distinction this essay has been at pains to draw. The missing arrest mechanism is precisely an instance of the closure or maintenance process the framework’s own vocabulary says must be separately identified: whatever halts the coarsening is functioning as the regime’s boundary-maintaining process, in the sense of the closure discussed in the section on structure and its regimes, and a physical account of a persistent pattern is incomplete until that process, and not merely the initial instability that produced the pattern’s first appearance, has been specified.

This is the sense in which persistent structure can be the product of instability rather than its casualty: what has to be regulated is not the presence of instability but its extent, and regulation of extent is itself a form of the same admissibility relation introduced in the previous section, now expressed as a constraint on which wavenumbers are allowed to grow. A system with no regulation at all, $\kappa = 0$, produces unbounded growth at arbitrarily short wavelengths and no persistent pattern whatsoever; a system with too much regulation suppresses the instability entirely and returns to the boring, homogeneous, structureless state. Persistent, finite-scale structure occupies the narrow regime between these two extremes, which is itself a nontrivial admissibility condition on the dynamics rather than a generic outcome to be expected from any unstable system.

None of this is offered as an empirical claim requiring immediate acceptance; it is offered as a demonstration that the mereological vocabulary of admissibility and regime-relative identity has a natural physical instantiation, and that the instantiation reproduces a known and independently motivated pattern, bounded spinodal instability generating persistent pattern, rather than inventing a new one for the occasion. Extending this sketch to cosmological scales, where the same qualitative logic has been proposed to apply to large-scale structure formation through entropy gradients and smoothing dynamics, remains a genuinely speculative extrapolation, and the essay flags it as such rather than presenting it as an established result; what is established, and what the extrapolation borrows, is only the general principle that regulated instability, not stability simpliciter, is what generates persistent, finite-scale structure in the first place.

6 Reconstruction after loss

A structure's history is not always available for direct inspection. Records are destroyed, intermediate states are unobserved, originals are lost, and the ordinary framing of this situation treats it as a deficiency: something that used to be knowable no longer is, and the honest response is to admit ignorance and move on. There is a more productive framing available, one that treats the loss of an original as an occasion for a different, more precisely specified kind of knowledge claim rather than as the simple absence of one.

What survives loss is frequently not the object itself but a set of constraints the object's prior existence imposed on everything downstream of it, and those constraints can be sufficient to specify a space of admissible histories even in the complete absence of the original. A destroyed manuscript can leave behind citations, quotations, and stylistic influences in later works, each of which constrains what the manuscript could have contained without determining it uniquely. A vanished species can leave behind fossilized traces, genetic signatures in surviving relatives, and effects on the ecosystems it once inhabited, each of which constrains its likely morphology and behavior without reconstructing it directly. An erased data file can leave behind fragments in a filesystem's unallocated space, references in other files that once pointed to it, and behavioral traces in the systems that depended on it, each of which constrains what the file must have contained without recovering it bit for bit.

Formally, let X_t be a present configuration and H a putative history leading to it. Reconstruction asks for the set of histories compatible with X_t and with whatever independent constraints are available,

$$\mathcal{H}(X_t) = \{H : H \text{ compatible with } X_t \text{ and with the available constraints}\},$$

and treats this set, not a single recovered history, as the object of knowledge. This reframes the goal of historical inquiry away from possession of an original and toward specification of an admissible space, and it is a strict improvement in situations where the original is provably unrecoverable

but the space of admissible histories is nonetheless narrow enough to be informative. The size of $\mathcal{H}(X_t)$ becomes the relevant measure of how much has genuinely been learned: a reconstruction project that narrows a space of a priori plausible histories from something enormous down to a small, well-characterized family has succeeded in an entirely meaningful sense, even though, and this is the point being pressed, it has not recovered the original and, in the relevant cases, cannot in principle do so.

This reframing also changes what counts as a legitimate objection to a reconstruction. The objection that a proposed history is merely one among several members of $\mathcal{H}(X_t)$ is not, by itself, a refutation of that history’s plausibility; it is simply an accurate description of the epistemic situation, which was never one of certainty and was never claimed to be. A refutation requires showing either that the proposed history lies outside $\mathcal{H}(X_t)$, that it is incompatible with the present configuration or the independent constraints, or that a competing history within $\mathcal{H}(X_t)$ is substantially better supported by whatever additional evidence distinguishes members of the set from one another. Much unproductive dispute about contested historical or archaeological reconstructions trades on conflating these two very different kinds of objection, treating mere membership in a nontrivial admissible set as though it were equivalent to having no evidential support at all.

This is worth stating as its own stage rather than folding it silently into repair, because it establishes a principle repair will need: that identity claims under loss are properly stated as membership in an admissible set rather than as certainty about a particular past state, and that this is not a concession to skepticism but the correct formal shape of the claim being made. The same logic, applied prospectively rather than retrospectively, is exactly what the continuation operator introduced later in this essay does for the future: it specifies not a single predicted successor state but an admissible set of them, and for exactly the same underlying reason, namely that the available constraints, whether they run from present traces backward or from a present state forward, typically underdetermine a unique outcome without thereby failing to constrain anything at all.

7 Repair: persistence surviving disruption

A structure that persists only so long as nothing disturbs it is a fragile achievement, and most real structures are disturbed constantly. A stronger notion of persistence includes disturbance and asks what happens next. Let a perturbation take a structure from an admissible state X to a damaged state X' , outside the previous admissible set. Several things can follow: dissolution, in which no further organization survives and the structure simply ceases; transformation, in which the system settles into a different regime with a different admissibility criterion and cannot meaningfully be said to be the same structure as before, even though something persists; or repair, in which some operator R carries X' to a state X'' that is once again admissible relative to the structure’s governing invariants,

$$X'' \sim_I X_{\text{prior}},$$

without requiring that X'' match the pre-damage state pointwise.

That last clause matters, and it is worth dwelling on because a demand for pointwise restoration is both too strong and beside the point. A wound heals into scar tissue rather than into the original dermal architecture, and scar tissue lacks hair follicles, sweat glands, and the original arrangement of collagen fibers, yet the healed skin restores the invariant that actually mattered, a continuous barrier maintaining the organism’s boundary against its environment, without restoring the microscopic arrangement that preceded the injury. A corrupted file is repaired into a functionally equivalent but bit-for-bit distinct replacement, restored from a backup that was itself written at a slightly earlier moment and therefore differs from the corrupted file’s own immediate predecessor

in ways no one will ever be able to specify exactly. A damaged institution reconstitutes its functions through new personnel occupying old roles, often with substantially different individual habits, styles, and even policies, and yet counts, both legally and in ordinary usage, as a continuation of the same institution rather than the founding of a new one. In every competently repaired system, what is restored is the relevant invariant, not the antecedent material configuration, and treating repair as failed whenever it does not reproduce the original state exactly is a category error about what repair is for, one that would make repair impossible in principle for any system whose material substrate turns over on a timescale shorter than the interval between injuries, which describes essentially every biological and institutional system of interest.

This raises the question of what licenses calling a given transformation a repair rather than merely another disturbance that happens to land somewhere convenient, a question that becomes urgent precisely because the definition of repair just given, restoration of an invariant rather than of a state, is permissive enough to be gamed. The answer requires what can be called a repair warrant: some independent basis, drawn from the system's own history, its governing constraints, or an external reference, for judging that X'' genuinely restores the relevant invariant rather than merely resembling a restoration. This is not a pedantic addition. Systems capable of modifying their own records are also capable of manufacturing an appearance of repair without the substance, and the difference between the two turns entirely on whether an independent warrant exists. A repaired database that also silently rewrote its own audit log has produced a resemblance of continuity without the warrant that would justify treating it as continuity in fact, since the very evidence that would distinguish genuine restoration from convenient fabrication has been destroyed by the same act that claims to be the restoration. The same distinction separates genuine institutional reform, in which restored function is externally verifiable against records the reforming body does not itself control, from institutional self-narration, in which only the account of restoration has changed while the underlying dysfunction, and the records that would reveal it, remain untouched or are themselves edited.

It is worth surveying, briefly, how the repair warrant is typically supplied across different regimes, since the mechanism varies considerably even though the underlying requirement, independence of the warrant from the thing being warranted, does not. In biological repair, the warrant is typically supplied by redundancy: multiple, partially independent copies or pathways exist, and damage to one can be detected and corrected by comparison against the others, as in DNA repair mechanisms that use one strand of a double helix as a template for correcting damage to its partner. In software repair, the warrant is typically supplied by an external specification, a test suite, a formal contract, or a separately maintained backup, against which a proposed fix can be checked without relying on the system's own, potentially compromised, self-report. In institutional repair, the warrant is typically supplied by external audit, independent oversight, or a documentary record maintained by a party without an interest in the outcome of the repair being judged successful. In each case, the warrant's independence is what does the actual epistemic work, and its absence is what separates a system that can genuinely recover from damage from one that can only ever produce a plausible-looking account of having done so.

Repair, so understood, also has scale-crossing character that mirrors the scale-crossing character of damage discussed in the earlier section on structure and regimes. A repair intervention performed at one scale frequently targets restoration of an invariant defined at a different, usually higher, scale: a surgeon operates at the tissue scale in order to restore function at the organismal scale; a programmer edits a single function in order to restore correct behavior at the scale of an entire application's user-facing behavior; a central bank adjusts a narrow technical parameter, an interest rate, in order to restore a broader macroeconomic invariant, price stability, that is not itself directly manipulable at the scale where the intervention actually occurs. This cross-scale character

is part of why repair is difficult to formalize in full generality: a warrant that is entirely convincing at the scale of the intervention, the tissue heals cleanly, the function passes its unit tests, the interest-rate change is implemented correctly, does not automatically transfer to a warrant at the scale where the invariant of actual interest is defined, and a substantial amount of what goes wrong in poorly designed repair processes across all of these domains traces back to exactly this gap between local and global warrant.

8 Continuation: history as a constraint on the future

A structure's present state, even together with a correct account of its recent repairs, still underdetermines what can happen to it next unless its full relevant history is taken into account. Two systems can occupy identical present states while differing in which future transformations remain available to them, because history narrows possibility even after it has stopped being directly visible in the present configuration. A patient who has recovered fully from a prior injury and a patient with no such history can present with identical current measurements while differing substantially in their prognosis under a subsequent, similar stress, precisely because the first patient's tissue has already been remodeled by the earlier repair in ways that constrain, without necessarily showing up in, its response to a second insult. This licenses treating a system not as a state X_t alone but as a pair (X_t, H_t) , and defining a continuation operator. Some care is needed in stating this formally, since the successor states reachable from different present states are not automatically comparable objects; the fix is to fix in advance a common space Ω of possible future trajectories or terminal outcomes, relative to which any given history's continuation set is a subset,

$$\mathcal{C}(H_t) \subseteq \Omega,$$

rather than treating $\mathcal{C}_{H_t}(X_t)$ as a free-floating set of successor states specific to X_t with no shared space against which two different histories' sets could be meaningfully compared. History does not determine a single point of Ω ; it restricts the reachable region, typically without narrowing it to a single point. This is the sense in which the earlier claim, that history is a constraint on what can coherently happen next rather than a mere record of what already happened, becomes precise, though the precise version is more guarded than the informal one: two histories are equivalent only relative to a specified horizon, state description, and family of interventions $\mathcal{I}$ against which their continuation sets are being compared,

$$H_1 \sim_{\Omega, \mathcal{I}} H_2 \iff \mathcal{C}_{\mathcal{I}}(H_1) = \mathcal{C}_{\mathcal{I}}(H_2),$$

and two histories that satisfy this relative to one choice of Ω and $\mathcal{I}$ need not satisfy it relative to another. The unqualified claim, that a history's operational content is exhausted by the shape of the continuation set it generates, has to be read with this relativization attached, since two histories can permit identical immediate futures under one family of interventions while differing in evidential warrant, in who is responsible for what, in latent fragility that only a different, more searching family of interventions would expose, or in how each would respond to some later piece of information neither history has yet encountered. Collapsing this distinction, treating any two histories with matching near-term continuation sets as simply the same, would erase exactly the kind of historical detail the repair-warrant discussion above insists must be tracked and preserved, since a repair warrant's independence is itself a fact about history that need not show up at all in the immediate shape of $\mathcal{C}(H_t)$.

A further correction is needed to the claim that history only ever narrows this set. Some commitments do narrow it: an event that adds a binding obligation without discharging any prior one

shrinks the reachable region,

$$\mathcal{C}(H_{t+1}) \subseteq \mathcal{C}(H_t),$$

and this narrowing case is the one most systems spend most of their time inside, since ordinary specification, decisions, promises, design choices, is overwhelmingly narrowing in exactly this sense. But narrowing is not the only thing history does, and treating it as though it were would misdescribe repair, learning, and institutional reform, each of which is capable of enlarging the reachable region relative to what a damaged or uninformed prior state permitted, restoring or even expanding possibilities that had been foreclosed rather than merely preserving whatever possibilities happened to remain. A course of physical therapy that restores a range of motion lost to injury, a repaired institution that regains a capacity a period of dysfunction had eliminated, a skill newly acquired that opens options unavailable before it was learned: each of these is a case of $\mathcal{C}(H_{t+1})$ properly containing possibilities absent from $\mathcal{C}(H_t)$, not a case of further narrowing, and the regeneration of foreclosed possibility, rather than the mere preservation of whatever possibility survives, is precisely one of the central subjects this essay's treatment of repair has been building toward. The general picture is therefore not monotonic narrowing punctuated by occasional rupture, but a continuation relation capable of narrowing under commitment, enlarging under repair or learning, and, in the rupture case, being replaced wholesale by an unrelated region with no principled relation to what came before; distinguishing genuine enlargement, which can be traced, warrant in hand, to an identifiable repair or an identifiable increment of learning, from a rupture that merely resembles enlargement is exactly the kind of judgment the repair warrant of the previous section was built to support. This distinction matters a great deal in systems that accumulate history over long interactions, conversations, ongoing collaborations, evolving institutions, legal proceedings that build a case incrementally out of successive admissible findings, because the health of such a system can often be read off from whether its continuation region is narrowing toward, or being restored toward, something still recognizably itself, as against being replaced wholesale by an unrelated region with no such relation. The first two are ordinary, healthy transformation, the way a promising early sketch narrows, through successive decisions, into a finished and internally consistent design, or the way a damaged capability is genuinely restored by a traceable repair. The third is a rupture dressed up as continuity, the way a change of leadership can be described in an institution's own communications as continuous evolution while in fact discarding every commitment the institution had previously made.

It is useful to introduce a further piece of vocabulary here, that of a persistent operator, to describe constraints that remain active across many transitions rather than applying only to the single step that introduced them. A system prompt governing an extended interaction, a physical conservation law governing a sequence of dynamical states, a constitutional provision governing a long sequence of institutional decisions, and a personal commitment governing a long sequence of individual choices are all, on this account, the same kind of formal object: an operator Π that acts not on a single transition but on the entire continuation relation,

$$\Pi : \mathcal{C}_t \rightarrow \mathcal{C}_{t+1},$$

filtering out, at every subsequent step, whichever candidate successor states would violate the constraint the operator encodes, regardless of how many steps have elapsed since the operator was first introduced. The central diagnostic question for any long-running system, on this account, is which of its operative constraints are genuinely persistent in this sense, remaining active and enforced at every subsequent step, and which are merely local, applying once and then quietly lapsing without anyone having decided to discharge them; a great deal of dysfunction in long-running human and institutional processes alike can be traced to constraints that were assumed to be persistent operators but were, in the relevant systems, only ever local ones.

A concrete and unusually clean illustration of this machinery is furnished by an extended conversational or agentic interaction, treated as a continuation system in its own right rather than merely as an application of the general framework. A conversation is not merely a sequence of turns, $p_1, p_2, p_3, \dots$; it is a recursively constrained history,

$$H_{t+1} = H_t \cup \{p_t, r_t\},$$

in which each new prompt and response pair modifies the space of admissible future responses, sometimes by adding a persistent operator, an instruction meant to govern every subsequent turn, and sometimes by adding only a local one, a request that applies to the immediate response and then lapses. A system that fails to distinguish these two kinds of addition to its history will either forget instructions that were meant to persist or, just as problematically, treat one-off requests as though they had been meant to govern everything that follows, and both failure modes are naturally described as failures of continuation-tracking rather than as failures of any single response taken in isolation.

This licenses one further, sharper distinction, between two questions that are easy to conflate but come apart under exactly the conditions where the difference matters most. The first question is whether a proposed next state is, considered on its own, locally acceptable: call this verification, a check performed against the present state in isolation, asking only whether X_{t+1} is true, correct, or otherwise satisfactory judged by criteria that do not look beyond X_{t+1} itself. The second question is whether admitting X_{t+1} into the history preserves the system’s capacity to remain coherent afterward: call this commitment, and note that it is a question about $\mathcal{C}_{H_{t+1}}(X_{t+1})$ rather than about X_{t+1} alone, asking whether the continuation set that results from accepting X_{t+1} still contains states the system can actually live inside, or whether it has been narrowed down to nothing viable. A locally plausible addition to a history can still be a poor commitment, if what it authorizes downstream turns out to be incompatible with everything the history has already committed to; a contract clause that is individually reasonable can render an entire agreement unworkable once combined with clauses agreed to earlier; a single architectural decision that is locally sensible can foreclose every subsequent option that would have let a software system scale, without that consequence being visible at the moment the decision was verified as correct. Systems that evaluate only the first question while ignoring the second tend to look competent at every individual step and incoherent in the aggregate, which is precisely the failure mode that a purely snapshot-based standard of correctness cannot detect and that a continuation-based standard is built to catch, and the gap between the two standards widens, rather than narrows, exactly in the high-stakes cases, long negotiations, extended engineering projects, sustained institutional policy, where getting only the first question right and neglecting the second does the most damage.

9 Cognitive and semantic continuation

The trajectory developed so far can be traced through cognitive and semantic systems in a way that is neither a loose analogy nor a claim of literal identity with the physical and mereological cases already discussed, but a genuine instance of the same formal problem posed at a different regime.

Consider first the question of cognitive adaptation. A cognitive system that learns is, definitionally, a system that changes its own internal state in response to experience, and the central difficulty this poses is exactly the difficulty posed by structure and repair above: how can a system modify itself, sometimes substantially, without thereby losing continuity with its own prior state, in the sense of ceasing to be recognizably the same cognitive agent pursuing anything like the same

commitments it held before the modification. An optimization process that adjusts a system's internal parameters in response to feedback, generated from the system's own recent behavior and applied continuously rather than only at isolated retraining events, is one candidate mechanism for achieving this, precisely because it modifies the system incrementally, from its own current state, rather than replacing that state wholesale with an unrelated one computed from scratch. The formal requirement this imposes is exactly the narrowing condition introduced in the previous section: adaptation should shrink or reshape the system's admissible future behaviors in a way that remains traceable, warrant in hand, back through the sequence of experiences that produced it, rather than substituting an unrelated policy that happens to perform better on some externally imposed metric while discarding everything the system had previously learned to value or attend to.

Memory plays a specific and somewhat narrower role in this picture than ordinary usage suggests. Memory is not, on this account, best understood primarily as a repository, a store of past states retrieved intact on demand, although that is one thing memory systems sometimes do. Memory's more fundamental role is as a constraint on future processing: whatever influence a past event continues to exert on present computation, whether or not that influence takes the form of an explicitly retrievable record, is memory in the sense relevant here, and can be written, generically, as a mapping from history to continuation set, $H_t \mapsto \mathcal{C}_{H_t}$. This formulation covers cases that a pure storage-and-retrieval account handles awkwardly: a skill that has become automatic no longer requires conscious retrieval of the experiences that produced it, and yet those experiences continue to constrain present behavior exactly as memory should; conversely, a system that can retrieve an accurate record of a past event but whose present processing is entirely unaffected by that record, in any way that shows up in its actual behavior, has memory only in the weakest, most inert sense, since the record exerts no continuation-shaping influence at all. Treating memory functionally, as a constraint on continuation, unifies these cases in a way that treating it as storage does not.

The semantic case introduces a further wrinkle worth isolating on its own. A concept or a meaning is a paradigm case of a structure whose persistence does not depend on symbolic identity: a concept can survive translation into an entirely different language, with no phonetic, orthographic, or etymological continuity whatsoever between the original term and its translation, provided the pattern of inferential relations the term participates in, what it implies, what it is implied by, what it licenses inferring in combination with other concepts, is adequately preserved by the translation. Symbolic identity is neither necessary nor sufficient for semantic persistence: a term can keep its exact surface form across a period of linguistic drift while its inferential role shifts so substantially that using it in an old text and a new one amounts, for practical purposes, to using two different concepts that happen to share a spelling, and a term can change its surface form entirely, through translation or through the ordinary evolution of technical vocabulary, while its inferential role remains essentially untouched. This is the semantic instance of the distinction, drawn earlier in the discussion of distinction generally, between the criterion that actually does the identity-preserving work and any more superficial marker, spatial contiguity in the physical case, surface form in the semantic case, that is neither necessary nor sufficient for it.

A closely related point concerns the maintenance of relevant context during interpretation, a task that is dynamic rather than static in exactly the sense the earlier discussion of informational boundaries anticipated. A system interpreting an ongoing sequence of input, whether a sentence, a conversation, or a longer discourse, must continuously decide what counts as relevant background for interpreting the next increment and what has become irrelevant and can be safely discarded or deprioritized, and this decision is not made once, at the outset, but is remade continuously as the discourse unfolds. This is, in a precise sense and not merely a suggestive one, the same operation a Markov blanket performs in regulating which variables belong to a system's interior

and which belong to its environment: the boundary between relevant and irrelevant context is not fixed in advance but is itself part of what the interpreting system dynamically maintains, narrowing and reshaping as new material arrives in a manner directly analogous to the narrowing of the continuation set $\mathcal{C}_{H_t}$ discussed above, since what counts as relevant context just is, functionally, a specification of which parts of the accumulated history are still doing constraining work on the interpretation of what comes next.

10 Computation as transformation ontology

The trajectory also motivates a specific, and somewhat unusual, way of thinking about what computation fundamentally is, one that treats programs as processes of growth, composition, transformation, and controlled collapse rather than as static manipulations of inert data structures by external operations.

The conventional picture of computation, inherited largely from the mathematics of functions applied to arguments, treats data as passive and operations as the active ingredient: a data structure sits still until some external function reaches in and transforms it, and the data structure itself has no persistence-relevant properties beyond its current, momentary content. This picture is not wrong, in the sense that it is not inconsistent with anything computers actually do, but it is a poor fit for describing systems, and there are many of practical importance, in which the entities being computed over are themselves best understood as persistent structures undergoing the same kind of admissible transformation discussed throughout this essay, rather than as inert values being pushed around by an entirely separate transformation apparatus.

An alternative computational ontology treats the basic unit of computation as a persistent, transformable structure that is seeded, grows by accumulating internal organization, composes with other such structures according to well-defined rules, transforms in response to its own internal state or to external input, and eventually collapses, either into a simpler resulting structure or into nothing, once its governing invariants can no longer be maintained. Schematically, the lifecycle runs

seed $\rightarrow$ growth $\rightarrow$ composition $\rightarrow$ transformation $\rightarrow$ collapse,

and the interesting design question, on this view, is not primarily what operations are available to act on data, but what invariants a given structure maintains across each stage of this lifecycle, and what repair mechanisms, in the specific technical sense developed above, are available when a structure's invariants are violated partway through. This is a more faithful description than the conventional picture provides for a wide range of systems that programmers already build in practice, including any system organized around persistent, mutable, but invariant-preserving objects, any system built around message-passing between long-lived, stateful components, and any system, biological simulations, physical simulations, long-running services, in which the entities being modeled are themselves naturally described as things that grow, interact, and eventually terminate rather than as static values awaiting an external transformation.

A further consequence of adopting this ontology concerns recursive, multiscale structure formation, which recurs across a wide range of computational systems, and which is naturally described as a process in which the state of a structure at one level of granularity is produced by an operator that jointly depends on the previous state at that level, on some noise or unmodeled variation, and on the accumulated trajectory that produced the previous state, rather than depending only on the previous state in isolation. Schematically,

$$S_{n+1} = \mathcal{T}(S_n, \eta_n, \tau_n),$$

where S_n is the structural state at step n , η_n represents noise or unmodeled variation entering at that step, and τ_n represents the trajectory, the relevant portion of the history, that produced S_n . This formulation makes explicit a feature that a purely Markovian description, one in which S_{n+1} depends only on S_n , would obscure: two structures with identical present state S_n can nonetheless differ in their admissible next states S_{n+1} , because they differ in the trajectory τ_n that produced the shared present state, and this trajectory-dependence is exactly the computational instance of the claim, made in the section on continuation above, that a present state alone underdetermines a system's future without its history also being specified.

11 Biology and longevity

The trajectory reframes a specific and consequential biological question, the question of aging, in a way that does not require committing in advance to any particular molecular mechanism, and that connects the phenomenon directly to the vocabulary of repair and continuation developed above rather than treating it as a separate topic requiring its own independent framework.

An organism persists despite substantial material turnover, though the turnover is considerably less uniform than a first pass at the point might suggest. Many proteins, lipids, and, on longer timescales, many of the cells composing several tissues are replaced repeatedly over the course of a normal lifespan. But this is not true of all molecular material, and the exceptions are not incidental: genomic DNA in most cell types is not wholesale replaced but maintained and repaired in place, some extracellular-matrix components turn over on timescales of years to decades, lens crystallins in the eye are essentially never replaced after early development and must last a lifetime, and neurons and cardiac muscle cells in several tissues are themselves long-lived, non-dividing cells whose own molecular contents turn over unevenly around a comparatively stable cellular scaffold. The more defensible statement is that an organism's molecular material turns over on markedly heterogeneous timescales, some components replaced within days, others essentially never replaced at all, so that an organism's identity cannot be a fact about which particular molecules currently compose it, without this requiring the stronger and less accurate claim that turnover is near-total or uniform across the organism. This is the biological instance of the general point made at the outset of this essay, that many of the things people are most inclined to treat as stable, persisting objects are in fact processes maintained by continuous, if unevenly distributed, replacement and repair of their own material, and it licenses treating an organism, formally, as a system whose persistence depends on the continued operation of repair mechanisms, DNA repair, protein turnover and quality control, cellular replacement where it occurs, immune surveillance, operating over exactly this heterogeneous landscape of components rather than uniformly across all of them.

On this framing, aging is not merely the accumulation of damage to otherwise fixed parts, since the parts subject to accumulation are, for many but not all components, being continuously replaced regardless of age. But it would be an overcorrection to replace a single-damage theory of aging with an equally single-factor theory pinned entirely on the decline of repair mechanisms. Aging more plausibly involves interactions among several distinct processes: accumulating damage in the long-lived components that are not replaced, declining or misdirected repair and quality-control capacity in the components that are, dysregulation of the signaling programs that coordinate repair, replacement, and cellular behavior across tissues, developmental programs that were adaptive earlier in life and become maladaptive once continued into later life, resource and energetic constraints that limit how much repair and replacement can be sustained simultaneously, and historically accumulated state changes, epigenetic or structural, that repair machinery may faithfully and successfully maintain rather than correct, since such machinery is generally built

to preserve a cell's current state rather than to judge whether that state itself has become a liability. Represent the organism's biological continuation capacity schematically as $\mathcal{C}_{\text{bio}}(t)$, the space of future physiological states reachable from the organism's current condition without loss of the invariants that define it as a living, functioning organism of its kind, and aging becomes the empirical claim that

$$\mathcal{C}_{\text{bio}}(t + \Delta t) \subset \mathcal{C}_{\text{bio}}(t)$$

for successive intervals Δt across an organism's later life: a genuine and progressive contraction of viable future organization, produced by the interaction of damage, declining or misdirected repair, regulatory drift, resource constraints, and accumulated state change together, rather than by any one of these in isolation. This reframing does not by itself explain why the contraction occurs, and it deliberately declines to commit to any single mechanism among cellular senescence, the erosion of epigenetic regulatory fidelity, the depletion of stem-cell reserves, or the other candidate processes currently under active investigation; the value of the reframing lies instead in the intervention target it identifies, which is the joint maintenance of several dimensions of physiological continuation capacity at once, rather than the elimination of any single category of damage considered in isolation. This is a considerably weaker and more defensible claim than the suggestion that repair capacity alone, on its own, could in principle absorb an arbitrarily large rate of damage indefinitely; no finite repair system has unbounded detection, correction, replacement, or energetic capacity, and a more accurate statement of the framework's own commitment is that a repair system can maintain an organism's governing invariants only so long as the rate and distribution of damage it faces remain within those finite capacities, a threshold condition rather than an open-ended guarantee.

This reframing suggests, as a research hypothesis rather than as a consequence the framework has already established, that interventions maintaining multiple dimensions of repair, replacement, and regulatory fidelity simultaneously may prove more robust across a population and across the life course than interventions eliminating one narrowly specified damage source in isolation, since the latter, even when fully successful at their stated target, leave the remaining dimensions of physiological continuation capacity, including the several distinct processes just enumerated, untouched. That hypothesis is stated at this deliberately guarded level, rather than as a categorical prediction that multidimensional interventions should outperform narrower ones in every case, because the actual comparison depends on which damage source, which intervention, which outcome measure, and which population are specified, and the framework developed here does not by itself settle any of those choices; what it supplies is a reason to expect the comparison to favor breadth of maintained capacity over narrowness of targeted damage, all else equal, not a guarantee that all else will in fact be equal in any given empirical test.

12 Delegation, recovery margin, and contestability

The trajectory extends, with a change of emphasis rather than a change of formal apparatus, to systems involving delegation, cases in which an agent, human or artificial, acts on behalf of another party, and to the further, population-scale question of what happens when many such delegations occur within a shared social or technological system.

The central question in the individual case is not simply whether a given delegation increases what is immediately achievable, since almost any competent delegate will do that; automating a task, hiring an assistant, or adopting a new tool nearly always expands what a person can accomplish in the short run, and evaluating delegation solely by this measure would make it look uniformly beneficial regardless of its other properties. The more discriminating question is whether

the delegation preserves the delegating party's own capacity to resume the activity independently, should that become necessary or desirable, a property distinct from and not automatically implied by the immediate capability gain. Introduce a recovery margin, $\rho(A, H)$, representing something like the distance between a person's current, agentically mediated state with respect to some activity A and a state from which that person could resume independent control of A given their accumulated history H of reliance on the delegate. A small recovery margin indicates that a person retains most of the skill, knowledge, or infrastructure needed to resume the activity on short notice; a recovery margin that grows steadily over an extended period of reliance indicates a form of dependency that the immediate capability measure, focused only on present output, does not detect at all, since output can remain high, or even improve, throughout the entire period in which recovery margin is silently eroding.

This licenses a specific normative proposal, framed here descriptively rather than as an unargued assertion, that good assistance should be evaluated as a joint condition rather than a single measure:

$$\text{good assistance} = \text{capability expansion} + \text{preserved reversibility}.$$

An assistant, human or artificial, that expands what a person can do while steadily eroding that person's capacity to do it unassisted is trading a visible, easily measured short-term gain for a much less visible, and typically much larger, long-term cost, and the framework developed in this essay supplies the vocabulary needed to notice the trade before it has fully compounded rather than only after the recovery margin has shrunk to the point where its erosion becomes obvious from outcomes alone. This is naturally paired with what can be called a right of first refusal, the principle that a delegate should not silently occupy a domain from which the delegating party can no longer practically return, an occupation that can happen gradually and without any single decision point at which it would have been natural to object, which is precisely what makes the recovery-margin framing useful: it supplies a way of tracking the relevant quantity continuously, rather than relying on the delegating party to notice, after the fact, that some threshold has already been crossed.

Scaling this concern from an individual delegating party to a population yields a related but distinct notion, contestability, which asks not whether any one person retains a recovery margin with respect to their own activities, but whether a system, a platform, a market, an institution, leaves meaningfully available alternatives for the population subject to it as a whole. A system can preserve every individual participant's narrow recovery margin, each person could, in principle, walk away from this particular service, while nonetheless being structured so that walking away brings no one to any genuinely different alternative, because every apparent alternative is either owned by the same party, dependent on the same underlying infrastructure, or has itself been rendered noncompetitive by the dominant system's own operation. Contestability, unlike individual recovery margin, is therefore a property of the space of options available across a population rather than a property of any one participant's relationship to a single delegate, and a system can score well on the individual measure while scoring poorly on the population measure, which is exactly the situation that arises when a dominant platform's terms of service are, individually, perfectly reversible for any single user while the aggregate structure of the market leaves no genuinely independent alternative for users collectively to move to.

This supports a compact, if necessarily informal, characterization of institutional power in the vocabulary developed throughout this essay: power, in the sense relevant to contestability, is control over the space of admissible continuations available to others,

$$\text{power} \sim \text{control over admissible continuation},$$

a formulation that connects institutional and political analysis directly back to the continuation operator $\mathcal{C}_H$ introduced earlier, treating a dominant system's ability to shrink the continuation set available to a population, foreclosing alternatives that would otherwise have remained admissible, as the operative content of the otherwise somewhat impressionistic claim that such a system has become powerful over that population, rather than treating power as an independent primitive requiring its own, unconnected theoretical apparatus.

13 The mathematical vocabulary of invariance

The concepts developed informally throughout this essay, admissible transformation, invariance up to a specified equivalence, regime-dependent identity, borrow, without needing to import wholesale, several existing mathematical languages that were independently developed to formalize closely related ideas, and a brief survey clarifies both what these languages contribute and what they do not, on their own, supply.

Symmetry, in the classical sense formalized by group theory, supplies the most direct precedent: an invariant I associated with a group G of transformations satisfies $I(gX) = I(X)$ for every $g \in G$, which is a special, exact case of the admissibility relation $\mathcal{A}(X)$ introduced above, one in which the admissible transformations form a group and the invariant is preserved exactly rather than only up to some coarser equivalence. The framework developed in this essay generalizes this classical notion in two directions at once: the admissible transformations need not form a group, since repair operators in particular are typically not invertible, a repaired system cannot generally be un-repaired back into its damaged state by any operation available within the framework, and the preserved invariant need not be exact, since the equivalence relevant to biological, institutional, and semantic persistence is almost always a coarser one, $X \sim_I Y$ for some specified family of admissible differences, rather than strict identity.

Topology contributes a further, related idea: that some properties of a structure survive continuous deformation even when no rigid symmetry is present to guarantee it, which is precisely the flexibility needed to describe a river's persistence through gradual, irreversible changes of course, or a melody's persistence through the continuous variation of tempo and dynamics a performer introduces in interpretation, cases in which no group of transformations relates the initial and final configurations exactly but a topological or near-topological notion of equivalence nonetheless captures the sense in which they remain, informally, the same. Category theory supplies a third and more general contribution, a vocabulary for composition and structure-preserving transformation that does not presuppose any particular notion of space or continuity at all: a morphism $f : X \rightarrow Y$ together with a further morphism $g : Y \rightarrow Z$ compose into $g \circ f : X \rightarrow Z$, and this composition is exactly the abstract shape of a history, a sequence of admissible transformations chaining together into a single overall transformation from an initial to a final state, which makes category theory a natural, if so far only partially developed, candidate language for formalizing the continuation operator $\mathcal{C}_H$ introduced earlier in terms of composable morphisms constrained by the admissibility relations appropriate to a given regime.

None of these mathematical languages, on their own, supplies the specific claims this essay has made about repair warrants, recovery margins, or the biological interpretation of aging; they supply, rather, a formal grammar within which claims of that kind can eventually be stated with the precision that would allow them to be checked, extended, or refuted, and part of what remains to be done, flagged explicitly in the open problems below, is the work of actually carrying out that formalization for the specific, regime-dependent claims made in the sections above rather than resting content with the informal but, it is hoped, sufficiently precise versions of them given there.

14 The trajectory across regimes: a synthesis

Having traced entropy, distinction, structure, persistence, repair, and continuation through physical, mereological, historical, cognitive, semantic, computational, biological, and institutional regimes in turn, it is worth stating explicitly what has and has not been established by doing so.

What has been established is that the same six-stage skeleton produces sensible, non-vacuous questions in every regime it has been applied to: a question about what resists dissolution, a question about what makes a candidate structure identifiable in the first place, a question about which transformations that structure can undergo without ceasing to be itself, a question about how it maintains that admissibility over time, a question about how it recovers when that maintenance locally fails, and a question about how its accumulated history shapes what can happen to it next. In every regime surveyed, physical field configurations, biological organisms, semantic concepts, computational structures, delegated agentic systems, and social institutions, this sequence of questions turns out to be worth asking, and asking it surfaces a formal structure, admissible sets, repair warrants, continuation operators, narrowing versus override, that recurs with enough specificity across regimes to be worth treating as one general framework rather than as six separately motivated vocabularies that happen to resemble one another superficially.

What has not been established, and what this essay has tried consistently not to overclaim, is that any one of these regimes reduces to any other, or that the physical sketch offered in an earlier section is anything more than a demonstration that the general vocabulary has at least one concrete instantiation with independently motivated technical content. The biological reframing of aging does not depend on the physical sketch being correct; the institutional analysis of contestability does not depend on the biological reframing being correct; each regime's application of the trajectory stands or falls on its own evidential and conceptual merits, and the unifying claim is limited to the observation that the same formal questions, asked regime by regime, produce a coherent and mutually reinforcing pattern of answers rather than a scattered and unrelated one.

A useful way of summarizing the whole picture is as a ladder of regimes,

physical → biological → cognitive → semantic → computational →
institutional → civilizational,

with two things flagged explicitly about this ladder rather than left implicit. First, it is not a claim of reduction: lower rungs are not more fundamentally real than higher ones in any sense this essay needs to defend, and nothing above has shown, or tried to show, that the higher rungs are nothing but complicated arrangements of the lower ones. Second, and more narrowly, the ladder as drawn corresponds to the regimes this essay has actually examined; a fuller ladder could in principle be extended downward through chemical, atomic, and quantum regimes, and upward past the civilizational, but this essay has not analyzed those further rungs independently and the ladder should be read as marking the range actually traversed above, together with a prospective claim, not yet checked here, that the same six-stage question would continue to be worth posing at ranks this survey did not reach, rather than as a report that it has already been posed and answered at every rank a complete ladder might include.

15 Persistence as a condition on truth

One consequence of the framework is worth isolating because it inverts a default assumption about the relation between truth and time. Truth is ordinarily treated as prior: something either holds or does not, and only afterward does the question arise of whether anyone can access, remember,

or verify it. The framework developed here suggests the dependency runs at least partly the other way. Evaluating whether a claim is true requires a distinction to be held fixed long enough to compare it against evidence, requires the evidence itself to have persisted in some accessible form, and requires the evaluating system to remain continuous enough across the comparison that the comparison counts as one act rather than two disconnected ones performed by unrelated successor states. None of this is available for free. It depends on structures, records, memories, institutions of citation and replication, that have to be built and actively maintained against the same entropic pressure discussed at the outset of this essay, and every one of the mechanisms surveyed in the intervening sections, repair warrants, continuation-tracking, recovery margins, contestable alternatives, turns out to have a direct analogue in the practices that keep a body of knowledge verifiable over time: a scientific result that cannot be independently repeated has, in effect, no repair warrant; a citation record that can be silently altered by the party whose claims it is meant to substantiate has lost exactly the independence a repair warrant requires; an institution of peer review that becomes captured by the very community whose claims it is meant to check has lost its contestability in precisely the sense developed in the section on delegation and institutions above.

This does not demote truth to a matter of convention or make it observer-relative in any troubling sense. It relocates the demand: before a claim can be assessed as true or false, the apparatus needed to perform the assessment has to persist long enough to perform it, and that apparatus is exactly as vulnerable to dissolution, damage, and needed repair as any other structure discussed in this essay. A single act of verification is not exempt from the six-stage trajectory that governs everything else surveyed here; it is simply a case of the trajectory playing out over a timescale, and at a level of abstraction, that makes its dependence on prior persistence easy to overlook, since the persistence in question, of records, of shared vocabulary, of institutions capable of checking one another's work, has typically already been secured by the time any individual act of verification takes place, making that verification look self-standing when it is not. Persistence, in this narrow but load-bearing sense, precedes verification, not as a metaphysical priority claim about what is more fundamental in the universe, but as a straightforward structural precondition on any system capable of performing verification at all, and a precondition, moreover, that can fail quietly, through the slow erosion of exactly the repair warrants and contestable alternatives discussed above, well before its failure becomes visible in the form of any individual claim being shown false.

16 Open problems

The framework leaves several questions genuinely open rather than merely unaddressed for space, and it is worth stating them with enough precision that future work, by this author or by others, could in principle resolve them rather than merely gesture in their direction.

The first concerns invariant selection: something has to determine which features of a system count as the invariants whose preservation defines persistence, and it is not yet clear whether that selection can be derived from a system's own dynamics, identifying the invariants a system's own regulatory processes actually track and defend, or must be supplied externally by whoever is doing the distinguishing, which would reintroduce an observer-dependence this essay has tried to keep at arm's length by insisting on the distinction, in the earlier section on structure and its regimes, between ontic and merely epistemic partitions. A satisfactory answer would specify a procedure for reading a system's relevant invariants off its own dynamics, rather than relying on an external analyst's judgment about which of a system's many properties happen to seem important.

The second concerns cross-scale propagation: an adequate account should specify, rather than merely gesture at, the conditions under which damage at one regime reliably disrupts admissibil-

ity at another, and the conditions under which it does not, since the essay has so far only asserted, without formalizing, that this propagation is probabilistic rather than deterministic and that repair frequently operates across scales in the manner described in the section on repair above. A quantitative account of cross-scale propagation would need to specify, for a given pair of coupled regimes, something like a transfer function relating the probability of higher-scale disruption to the magnitude and kind of lower-scale damage, and no such account has been offered here beyond the qualitative observation that the coupling exists and is not one-to-one.

The third concerns whether the size or richness of a continuation set can be given a workable quantitative measure, since a usable measure of that kind, some functional $\mu(C_H(X))$ satisfying reasonable consistency conditions, would immediately supply operational definitions of resilience, autonomy, and recoverability that are currently only qualitative in this essay's treatment. It is not obvious that any single scalar measure could do justice to the relevant notion across all the regimes surveyed above, since the continuation set relevant to a biological organism's resilience and the continuation set relevant to an institution's contestability are unlikely to share enough formal structure to support a common measure without considerable additional work, and this essay does not attempt that work here.

The fourth, and least tractable, concerns empirical discrimination: which of the claims made along the way, particularly the physical sketch in the earlier section on localized persistence and the biological reframing of aging, generate predictions distinguishable from those of existing, independently successful theories, as opposed to redescribing the same predictions in new vocabulary. An honest accounting of the program's current state places most of its physical content on the wrong side of that line for now, a redescription rather than a novel prediction, which is a reason for continued scrutiny rather than a reason for abandoning the vocabulary, provided the vocabulary continues to earn its keep in the domains, mereology, repair, continuation, delegation, where it is doing genuine explanatory work independent of any specific field-theoretic commitment, as the sections on structure, repair, continuation, and delegation above have attempted to show it does. The biological reframing of aging fares somewhat better on this count, since it does generate a discriminable prediction, that interventions targeting general repair and replacement capacity should outperform interventions targeting any single category of molecular damage, even though that prediction remains to be tested with the rigor the framework's own standards would demand of it.

17 Conclusion

The claim advanced here is modest in one sense and ambitious in another. It is modest because it does not assert that biology, physics, computation, and institutional theory are secretly one subject matter; each retains its own content, its own methods, and its own standards of evidence, and nothing in this essay's survey across physical, mereological, historical, cognitive, semantic, computational, biological, and institutional material, categories chosen for expository convenience rather than as a settled taxonomy with a fixed count, has required treating any one of them as reducible to, or as merely an application of, any other. It is ambitious because it asserts that a single formal trajectory, from a background of entropic dissolution, through the establishment of a distinction, through the constraint on transformation that makes a distinction worth calling a structure, through the temporally extended survival of that structure, through its recovery from damage under an independent warrant, to the shaping of its future possibilities by its own history, recurs across these subject matters with enough specificity to be worth naming and studying as its own topic, and that naming it makes visible connections, between the biology of aging and

the vocabulary of repair warrants, between institutional power and the shrinking of continuation sets, between the persistence of a verified truth and the same repair-and-contestability apparatus that keeps any other structure surveyed here from quietly dissolving, that would otherwise remain scattered across disciplines with no shared vocabulary for noticing them.

What persists is never simply what a system is made of. It is the organization that keeps managing to remain itself while everything it is made of keeps being replaced, repaired, renegotiated, and rewritten. Structure through change is the name for that organization, and for the shared, cross-domain question of what has to hold for it to continue: not what stays the same, since in the interesting cases surveyed throughout this essay almost nothing does, but what keeps being regenerated well enough, and warranted independently enough, that the regeneration counts as the same thing continuing rather than something else quietly taking its place.