The Uses of the Unfinished:

Continuation, Foreclosure, and the Architecture of Commitment

Flyxion
Independent Researcher

September 27, 2026

Abstract

Completion is ordinarily treated as the successful terminal condition of production, and incompletion as a privation measured by distance from the object it has failed to become. This essay develops a contrary account built on one definition: completion is an authority's withdrawal of recognition from the continuations of an artifact within a declared scope. Completion is then not a property of a state but a relation among a state, an authority, a time, and a scope, and it is frequently a transfer of unresolved costs from an authorized producer to less powerful successors. The argument distinguishes *inert* incompletion, in which work has stopped without an intelligible means of resumption, from *latent* incompletion, in which the present state remains legible, accessible, and continuable along more than one admissible path, and it distinguishes three ways in which continuations are withdrawn: epistemically, operationally, and jurisdictionally.

A core model represents an artifact by its state space, recognized continuations, interpretive information, costs, allocation of authority, and identity relation. Descriptive continuational breadth is kept separate from signed, control-indexed option value, so that the value of a retained option depends on who exercises it and for whom. The principal results show that completion can raise present fitness while contracting breadth; that an option can be valuable to its controller and harmful to those it affects; that optimal closure over a dependency structure, condensed where commitments are mutually dependent, is a maximum-weight ideal that in general leaves purpose-like components open; that suggestion systems with sufficient refusal cost make arbitrary predictions self-confirming; that revision capacity can grow while the alignment of authority with affectedness declines; and that symmetric reopening rules erode settlements under repeated challenge. Standard results from the theory of irreversible investment and the value of information are identified as such. Domain models for development, interfaces, machine assistance, administration, political authority, and ruins instantiate the core model. The normative conclusion is a principle of *asymmetric completion*: close the transformations whose openness would destroy the conditions of participation, and preserve those through which future participants may revise purposes that no present authority can legitimately settle in advance.

Keywords: completion; option value; irreversibility; continuation; modularity; choice architecture; classification; institutional design.

Contents

1 Completion and Its Presumed Priority

An unfinished thing is usually interpreted through the completed thing it has failed to become. Exposed reinforcing bars anticipate another storey. A blank chapter condemns a manuscript by referring the reader to the absent chapter that ought to occupy its place. A temporary patch is understood not through the work it presently performs but through the proper repair for which it waits. A protocol without an implementation, a building without cladding, a road terminating in gravel, and an institution operating under provisional rules appear to share a subordinate ontological condition. Each is described as a deficient version of an intelligible future object. The completed object supplies the measure, and the unfinished object receives its identity as a distance from that measure. Even when the future object never arrives, it governs the interpretation of what exists.

This ordinary interpretation contains a concealed ordering relation. Let X denote an artifact and let x_t denote its state at time t. Suppose an observer possesses a terminal specification g, the state the artifact is intended to reach. The incompletion of x_t is represented by a distance

$$d(x_t, g) \;\geq\; 0, \tag{1}$$

where $d(x_t, g) = 0$ signifies completion and larger values signify greater deficiency. The metric may concern missing physical components, unsatisfied requirements, incomplete arguments, unimplemented functions, unresolved cases, or any combination of these. What matters is that the meaning of x_t is derived from g. The present state is not evaluated primarily by what it permits, supports, or preserves. It is evaluated by what remains to be done before the terminal specification is satisfied.

The model seems innocent because many tasks genuinely possess identifiable terminal conditions. A bridge must connect its banks and sustain specified loads. A proof must establish its conclusion. A medicine must have a determinate formulation before it can be administered responsibly. Yet the model incorporates two assumptions that cease to be innocent when generalized. The first is that the terminal state g is uniquely specified, or at least uniquely authoritative. The second is that movement toward g is monotonically beneficial, so that a reduction in $d(x_t, g)$ constitutes an improvement independently of the alternatives removed by that movement. Together these assumptions convert one possible continuation into the criterion by which all present continuations are judged.

1.1 Continuation sets and continuational breadth

To make the suppressed alternatives visible, let Ω be a state space containing the configurations the artifact might occupy, and let $A(x) \subseteq \Omega$ be the set of states reachable from x through admissible transformations. Admissibility is not identical to physical possibility. A wall can physically be demolished, a database can physically be replaced, and a statute can physically be repealed, but these transformations may be financially prohibitive, institutionally forbidden, informationally inaccessible, or destructive of the very object one hopes to revise. The continuation set $A(x)$ contains only those states reachable under the material, legal, epistemic, and organizational constraints that actually obtain. Section 2 makes these constraints explicit; here it suffices that $A(x)$ is a measurable subset of Ω.

The conventional account evaluates x principally through its distance from g. A continuational account asks instead about the structure of $A(x)$. The cardinality $|A(x)|$ is a crude measure of how many states remain accessible, and it is insufficient, because a large set of trivial variations may matter less than a small set containing substantively different futures. Equip Ω with a measurable, symmetric, nonnegative *distinction function* $\delta : \Omega \times \Omega \to [0, \infty)$ with $\delta(y, y) = 0$, where larger values indicate more consequential differences between states, and fix a reference measure ν on Ω.

Definition 1 (Continuational breadth)**.** The continuational breadth of a state x is

$$B(x) \;=\; \int_{A(x)} \int_{A(x)} \delta(y, z)\, d\nu(y)\, d\nu(z). \tag{2}$$

Two features of the definition should be made explicit. First, δ is not an objective distance. It is a declared distinction regime, formally a pseudometric chosen for the inquiry, and much of the normative work of the framework is done in declaring it: a room arrangement can be physically distant from another yet politically inconsequential, while two visually identical records can confer very different legal standing. Continuational breadth is always breadth relative to a declared distinction regime. Second, B combines two properties that are worth separating. Writing $\mathsf{m}(x) = \nu(A(x))$ for the *continuation mass* and, whenever $0 < \mathsf{m}(x) < \infty$,

$$\bar{\delta}(x) \;=\; \frac{1}{\mathsf{m}(x)^2} \int_{A(x)} \int_{A(x)} \delta(y, z)\, d\nu(y)\, d\nu(z) \tag{3}$$

for the *mean differentiation*, one has $B(x) = \mathsf{m}(x)^2\, \bar{\delta}(x)$. The decomposition distinguishes the loss of many similar continuations, which lowers m, from the loss of a few radically different ones, which lowers $\bar{\delta}$. It also shows that breadth depends on the scale of the reference measure, since replacing ν by $c\nu$ multiplies B by c^2. Comparisons of breadth are therefore meaningful only under a fixed ν, and the choice of ν is itself normative: a measure assigns weight to possible futures, and in an institutional setting it may encode whose futures count. Breadth is also purely descriptive. It measures how differentiated the reachable future is, not whether that future is desirable, so a system with more dangerous failure modes has, other things equal, a larger B. Nothing in the argument treats breadth as a good in itself. Valuation enters separately in Section 3, through signed, agent-indexed utilities computed over admissible continuations only.

The quantity $B(x)$ does not count mere motion. It estimates the *differentiated* future retained by the present configuration: it vanishes when $A(x)$ is a ν-null set or when all reachable states are mutually indistinguishable, and it is monotone under inclusion, since $A(x') \subseteq A(x)$ implies $B(x') \leq B(x)$ because the integrand is nonnegative. An artifact with a thousand cosmetic settings but a single permitted institutional purpose may therefore have high superficial variability and low continuational breadth, while an artifact with only a few attachment points may have greater breadth if those points support transformations of ownership, use, interpretation, or governance.

Completion can now be represented as a transformation $C : x \mapsto x'$ whose image satisfies a terminal predicate $G(x')$. Nothing in this definition entails $B(x') \geq B(x)$, and completion frequently produces the opposite inequality. A permanent wall may improve the present fitness of a dwelling while eliminating several future spatial arrangements. A published technical standard may permit interoperable implementation while making incompatible conceptual distinctions increasingly costly to introduce. A fixed database schema may make current queries reliable while excluding kinds of evidence unanticipated when its fields were chosen. A policy may resolve an immediate dispute while narrowing the vocabulary through which later claimants can become administratively visible. In each case completion adds a determinate capacity while removing a plurality of indeterminate ones.

Let $F(x, e_t)$ denote the fitness of state x relative to the environment e_t.

Proposition 1. *A completion transformation may increase the present fitness of an artifact while strictly decreasing its continuational breadth.*

Proof. Let $C : x \mapsto x'$ satisfy $F(x', e_t) > F(x, e_t)$. Suppose C fixes some component q which previously admitted values in a set Q, replacing Q with a singleton $\{q^*\}$, and suppose no other admissible transformations are created, so that $A(x') \subseteq A(x)$. Let $Y_q \subseteq A(x)$ be the set of states reachable only through a value $q \neq q^*$, and suppose $\nu(Y_q) > 0$ and $\nu(A(x')) > 0$, with $\delta(y, z) > 0$ for $y \in Y_q$, $z \in A(x')$ (the lost branches are substantively distinct from the retained ones). If no admissible transformation from x' restores Y_q at comparable cost, then $Y_q \cap A(x') = \emptyset$, and

$$B(x) - B(x') \;\geq\; 2 \int_{Y_q} \int_{A(x')} \delta(y, z)\, d\nu(y)\, d\nu(z) \;>\; 0.$$

Hence fitness rises while breadth strictly falls. $\square$

The proposition is elementary, but its implications are obscured by the language of progress. If completion is represented only as a decrease in $d(x, g)$, the contraction of $A(x)$ disappears from view. The object has moved closer to its goal, and the branches lost during that movement are classified as irrelevant because they do not terminate at g. The conclusion has been embedded in the metric: completion appears unambiguously productive because the accounting system recognizes the acquired state but not the extinguished option set.

1.2 Present fitness and retained continuability

A more adequate evaluation must consider both realized fitness and retained continuability. Let $\vartheta \in [0, 1]$ discount future adaptive value, and let $H(x) \geq 0$ denote the present cost of maintaining the artifact's openness. Define

$$V(x) \;=\; F(x, e_t) \;+\; \vartheta\, \mathbb{E}\Big[\max_{y \in A(x)} F(y, e_{t+1})\Big] \;-\; H(x). \tag{4}$$

The second term is the expected value of being able to select a later state after more information about the future environment e_{t+1} has become available. It is analogous to real option value, but the relevant uncertainty is not exclusively economic. It may concern future needs, newly recognized persons, revised scientific distinctions, altered ecological conditions, emergent hazards, or purposes not yet articulable within the existing institutional vocabulary.

Under (4), completion $x' = C(x)$ is justified only when its gain in present fitness exceeds both the adaptive value it forecloses and the costs it displaces:

$$F(x', e_t) - F(x, e_t) \;>\; \vartheta\left(\mathbb{E}\Big[\max_{y \in A(x)} F(y, e_{t+1})\Big] - \mathbb{E}\Big[\max_{z \in A(x')} F(z, e_{t+1})\Big]\right) + \big(H(x') - H(x)\big) + \Delta T, \tag{5}$$

where ΔT measures costs displaced to successors rather than eliminated. The first two terms on the right follow directly from requiring $V(x') > V(x)$ in (4); the maintenance term $H(x') - H(x)$ is negative when completion reduces upkeep, which is the precise sense in which saved maintenance can justify closure. When $A(x') \subseteq A(x)$ the bracketed option term is nonnegative, so a completion that saves maintenance must still pay for the options it destroys. The inequality is more demanding than ordinary project criteria because it refuses to identify the producer's closure with the artifact's total success. A building can be complete for the developer while remaining costly to alter for its occupants. A software release can be complete for the development schedule while transferring unresolved incompatibilities to users and support workers. A statute can be complete as legislation while producing appeals, exceptions, discretionary workarounds, and private injuries that never appear in the legislative account. The declaration of completion often marks the moment at which one agent ceases to bear responsibility for indeterminacy and others begin to bear its consequences.

1.3 The declaration boundary

Let the relevant agents be indexed by $i \in I$, and let $k_i(x \to y) \geq 0$ denote the cost borne by agent i when the artifact is transformed from x to y. Let $\mathsf{P} \subset I$ be the producers authorized to declare completion, and $\mathsf{S} = I \setminus \mathsf{P}$ the successors, users, maintainers, and affected parties. The producer-side cost of a completion C is

$$K_{\mathsf{P}}(C) \;=\; \sum_{i \in \mathsf{P}} k_i(x \to x'), \tag{6}$$

whereas its social cost is

$$K_{\text{all}}(C) \;=\; K_{\mathsf{P}}(C) \;+\; \sum_{j \in \mathsf{S}} \mathbb{E}\left[\min_{y \in A(x') \cap G_j} k_j(x' \to y)\right], \tag{7}$$

where $G_j \subseteq \Omega$ is the set of states that successor j will in fact need to reach, and the minimum over an empty set is taken to be the cost of replacement. An evaluation confined to K_{P} will systematically favor terminal forms that are inexpensive to deliver and expensive to inhabit, maintain, reinterpret, or reopen. The omission is not an accidental measurement error. It follows from the institutional boundary that determines when the project ends and whose work counts as part of it.

Observation 1. *If the authority that declares completion is evaluated only on costs incurred before declaration, it can improve its measured performance by transferring unresolved transformation costs beyond the declaration boundary.*

Proof. Suppose the producer chooses between C_1, which preserves an accessible attachment point at present cost $a > 0$ and leaves successors with expected adaptation cost 0, and C_2, which omits the attachment point at no present cost and imposes expected successor adaptation cost $b > a$. Then $K_{\mathsf{P}}(C_2) = 0 < a = K_{\mathsf{P}}(C_1)$, while $K_{\mathrm{all}}(C_2) = b > a = K_{\mathrm{all}}(C_1)$. An evaluator minimizing K_{P} selects C_2. Measured performance improves while total expected cost rises by $b - a$. □

This explains why apparently finished systems so often accumulate unofficial margins. Residents enclose balconies, cut paths across lawns, convert garages into bedrooms, and rearrange spaces designated for a single function. Workers construct spreadsheets beside enterprise systems because the official schema cannot represent the distinctions required by actual work. Users write scripts around applications that advertise completeness. Readers annotate printed arguments whose fixed surfaces cannot respond to later objections. These practices are usually described as misuse, workaround, customization, or technical debt. They are better understood as *continuational labor*: they restore branches that completion eliminated, or they construct attachment points where the official artifact supplied none. The quantity b in Observation 1 is, in practice, paid in exactly this labor.

The resulting picture reverses the presumed priority of completion. Incompletion is not necessarily a measurable distance from a privileged terminal object. Completion is one operation within a larger field of continuations, and its value depends on the capacities it realizes, the alternatives it extinguishes, the costs it redistributes, and the agents authorized to decide that no further recognized continuation belongs to the object. We may therefore state the working definition that the rest of the essay refines.

Definition 2 (Completion)**.** Let $\mathfrak{D}$ be a declared *scope*: a responsibility domain specifying which transformations of the artifact are at issue, whether structural, functional, semantic, jurisdictional, or some combination. The predicate $\mathrm{Complete}(x_t, \mathsf{P}, t, \mathfrak{D})$ holds when the authority P withdraws recognition from every transformation of x_t within $\mathfrak{D}$ other than those admitted by a stated maintenance or reopening procedure, maintenance being understood under a declared equivalence relation. Completion is thus not a property of a state but a relation among a state, an authority, a time, and a scope.

The scope parameter is not a technicality. Authorities rarely withdraw recognition from every continuation of an artifact. A completed software release still admits patches, a completed building admits renovation, a decided case admits appeal. What is completed is usually a phase, a version, a claim, or an assignment of responsibility. Scope also sharpens the declaration boundary: the boundary in Observation 1 is the boundary of $\mathfrak{D}$, and a producer can displace costs precisely by choosing a scope that excludes the transformations successors will need. An artifact can be complete within one scope and open within another, which is why the scope belongs in every declaration.

The central question is consequently not whether an artifact is finished but whether the closure it embodies is proportionate to the knowledge and authority available at the time of closure. The unfinished thing can still be defective. It may be dangerous, unintelligible, or abandoned, and nothing in the formalism converts absence into virtue. The formalism does, however, prevent closure from functioning as its own justification. A polished terminal state is not superior merely because its contingency has become difficult to see. An exposed joint may signify neglect; it may also preserve an intelligible site at which a future participant can enter. Distinguishing these cases requires a more exact taxonomy of incompletion.

1.4 The core model and the plan of the essay

Definition 2 is the essay's primary contribution, and the remainder is organized as its consequences. Continuational breadth describes what an authority withdraws when it completes. The declaration boundary describes who escapes the resulting costs. Option value describes what the withdrawal loses under uncertainty. Dependency ideals describe which withdrawals can nonetheless be justified. Continuation rights describe who may reopen what has been closed. To keep these consequences within one frame, every artifact studied below is treated as an instance of a *continuation structure*

$$\mathcal{C} \;=\; (\Omega,\; A,\; \Sigma,\; \mathcal{K},\; \mathcal{W},\; \sim_{\mathcal{R}}\,), \tag{8}$$

where Ω is the state space, A the recognized continuation relation, Σ the information required to interpret states, $\mathcal{K}$ the cost structure of transformations, $\mathcal{W}$ the allocation of authority to perform them, and $\sim_{\mathcal{R}}$ the equivalence relation under which the artifact's identity is preserved.

Sections 1 through 4 develop the general theory: breadth and the declaration boundary (Section 1), the distinction between inert and latent incompletion (Section 2), option value (Section 3), and the selection of closures over a dependency structure (Section 4). Sections 5 through 11 are domain models rather than extensions of a single calculus. Each instantiates (8) with its own state, continuation, information, cost, authority, and identity variables, stated at the start of the section, and each uses a deliberately stylized model to expose one mechanism. Together they demonstrate the portability of the core definitions; they do not constitute a unified quantitative theory of unfinishedness, and they are not offered as one. Section 12 assembles a decision rule, Section 13 takes up objections and boundary conditions, and the final section concludes. Standard results are identified as such where they appear, and results that follow directly from a chosen specification are labeled observations rather than propositions. The conclusion relates the framework to a companion account, admissible degradation, which asks what must survive when capacity contracts rather than what must remain open when a project declares success.

2 Inert and Latent Incompletion

Not every unfinished artifact preserves possibility. Some unfinished buildings admit rain until reinforcement corrodes and floors become unsafe. Some provisional institutions persist without acquiring either legitimacy or an orderly procedure for revision. Some abandoned programs contain thousands of lines of source code but no recoverable account of their invariants, dependencies, or intended behavior. In such cases the absence of closure does not preserve a useful future; it consumes one. Time acts on the unfinished artifact, degrading the materials, knowledge, trust, and coordination on which any later continuation would depend. An adequate theory must therefore distinguish an unfinished state that retains accessible alternatives from one that merely has not reached an announced terminus.

2.1 The situated artifact

Represent an artifact at time t by the tuple

$$\mathcal{X}_t \;=\; (x_t,\; \Sigma_t,\; \mathcal{T}_t,\; \mathcal{K}_t,\; \mathcal{W}_t), \tag{9}$$

where x_t is the observable state, Σ_t is the information available for interpreting that state, $\mathcal{T}_t$ is the set of technically possible transformations, $\mathcal{K}_t$ assigns agent-relative costs to those transformations, and $\mathcal{W}_t$ records the agents authorized or practically able to perform them. The representation separates properties usually collapsed into the single judgment that something remains unfinished. An artifact may be materially alterable but epistemically opaque. It may be intelligible but legally closed. It may admit a physically simple modification that no relevant agent can afford. It may preserve many possible transformations, all of which destroy the capacities that made it worth continuing. Physical incompletion is therefore neither necessary nor sufficient for practical openness.

Identity preservation is specified by a declared set $\mathcal{R}$ of retained properties, which may include physical fabric, recorded evidence, functional capacity, institutional memory, semantic identity, or relations of trust, together with the induced equivalence $\sim_{\mathcal{R}}$: $x \sim_{\mathcal{R}} y$ when x and y agree on every property in $\mathcal{R}$. For an agent i, an epistemic tolerance τ, and a resource bound κ, define the agent-relative continuation set

$$A_i(x;\tau,\kappa,\mathcal{R}) = \left\{ y \in \Omega \;:\; \begin{array}{l} \exists \text{ a path } x \to y \text{ in } \mathcal{T}_t \text{ performable by } i \text{ under } \mathcal{W}_t, \\ \text{interpretable from } \Sigma_t \text{ with error at most } \tau, \\ \text{of cost to } i \text{ at most } \kappa \text{ under } \mathcal{K}_t, \\ \text{and with } y \sim_{\mathcal{R}} x \end{array} \right\}. \tag{10}$$

The last condition is essential. If a program can be altered only by rewriting it from the beginning, or a building adapted only by demolition, change remains possible in a trivial physical sense, but the artifact does not supply a continuation. It supplies material, land, or information for a replacement. Continuation requires that significant structure survive the transformation. Write $B_i(x)$ for the breadth (2) computed over A_i.

Because different artifacts preserve different organization, $\mathcal{R}$ cannot be universal. A repaired bridge may preserve little original material while preserving its load path, location, and public function. A scientific archive may preserve its files while losing provenance, rendering its physical survival epistemically inadequate. A constitution may retain its text while changes in enforcement destroy the continuity the text nominally represents. Preservation is always preservation under a declared equivalence relation, and the choice of $\mathcal{R}$ must be stated, defended, and made available to criticism. Where the retained properties admit a defensible aggregation rule, the relation may be relaxed to a graded one: a function $P(x,y) \in [0,1]$ measuring the proportion of relevant organization preserved, with continuation requiring $P(x,y) \geq r_{\min}$. The graded form is optional. Fabric, function, provenance, memory, and trust are not naturally commensurable, and a scalar P should be used only where a particular inquiry supplies the weights. The requirement prevents "continuability" from becoming an ornamental synonym for change.

2.2 Two forms of incompletion

Let G be the completion predicate and $\varepsilon > 0$ a materiality threshold.

Definition 3 (Inert incompletion)**.** An artifact is *inertly incomplete* for agent i at time t when $\neg G(x_t)$ and $B_i(x_t; \tau, \kappa, \mathcal{R}) \leq \varepsilon$.

Definition 4 (Latent incompletion)**.** An artifact is *latently incomplete* for agent i at time t when $\neg G(x_t)$ and $B_i(x_t; \tau, \kappa, \mathcal{R}) > \varepsilon$.

The distinction does not depend on whether the artifact looks unfinished. A rough concrete structure is inert if its load paths are undocumented, its materials have degraded, and no agent has the authority or resources to intervene. A visually finished building can remain latently incomplete if its service cavities are accessible, its partitions alterable, and its structural system permits later reconfiguration. Habraken's distinction between support and infill treats the built environment as a hierarchy of control rather than a terminal composition: a stable support can be highly finished at one level while retaining incompletion at another, allowing occupants rather than distant producers to determine internal arrangements [13]. Brand's account of buildings adapting through layers that change at different rates similarly shows that endurance often depends on parts remaining independently revisable rather than on the structure achieving a final form [7].

Latent incompletion is therefore structured rather than absolute. It requires closure at some levels so that openness remains usable at others. A door is useful because its frame is stable. A programming interface supports unanticipated implementations because its syntax is determinate. A scientific disagreement remains productively open only if evidence, methods, and disputed interpretations are preserved precisely enough for later investigators to reconstruct it. Total fluidity does not maximize continuability, because a state that cannot be identified cannot serve as the origin of a reliable transformation.

2.3 The continuability functional

Let $L(x) \in [0,1]$ denote the legibility of the present state, $E(x) \in [0,1]$ the accessibility of its intervention points, and $D(x) \geq 0$ the substantive diversity of the states reachable through them (for instance $D = B$ normalized). Adopt the Cobb–Douglas specification

$$\Gamma(x) \;=\; L(x)^{\alpha_L}\, E(x)^{\alpha_E}\, D(x)^{\alpha_D}, \qquad \alpha_L, \alpha_E, \alpha_D > 0. \tag{11}$$

The multiplicative form expresses complementarity: if any factor vanishes, $\Gamma = 0$. It is one useful specification rather than a uniquely motivated one; in logarithms it imposes fixed elasticities and a particular form of substitutability among the factors, and the arguments below rely only on its complementarity. A perfectly documented artifact with no accessible intervention point cannot be continued by the relevant agent. An accessible artifact whose state cannot be interpreted invites alteration without responsible continuation. A legible and accessible artifact that permits only one predetermined successor is *deferred*, not open: its future has been delayed rather than pluralized.

Proposition 2. *Unconstrained malleability does not maximize continuability.*

Proof. Index designs by a malleability parameter $m \geq 0$, with L, E, D differentiable and positive. Suppose $D'(m) > 0$ throughout, $E'(m) \geq 0$ and bounded, and that beyond some m_0 components no longer remain stable long enough to establish dependable relations, so that $L'(m) < 0$ for $m > m_0$. Then

$$\frac{\Gamma'(m)}{\Gamma(m)} \;=\; \alpha_L \frac{L'(m)}{L(m)} + \alpha_E \frac{E'(m)}{E(m)} + \alpha_D \frac{D'(m)}{D(m)}.$$

If the elasticity of legibility loss dominates, that is, if $\alpha_L |L'/L| > \alpha_E E'/E + \alpha_D D'/D$ on some interval (m_1, ∞) with $m_1 \geq m_0$, then Γ is strictly decreasing there. In particular, if $L(m) \to 0$ as $m \to \infty$ while D and E grow at most polynomially and L decays exponentially, $\Gamma(m) \to 0$. Hence the supremum of Γ is not attained at maximal malleability. □

The result identifies a central error in celebrations of flexibility. Flexibility is frequently measured by the number of configurations an artifact can assume, without asking whether transitions among them remain intelligible, affordable, and identity preserving. A system in which every component can change at once may possess a large formal state space but little practical capacity for directed continuation. The user cannot isolate causes, predict consequences, or preserve functioning relations while intervening. Such a system is mutable but not governably unfinished.

2.4 Nominal and effective openness

The distinction between formal possibility and admissible continuation explains why source availability, transparency, or legal permission cannot by themselves establish openness. Let $T(x)$ be the set of states reachable by transformations not prohibited by formal rules. Ordinarily $A_i(x) \subseteq T(x)$, and the inclusion may be strict. A repository may license modification while requiring tacit knowledge held by a vanished team. A law may grant a right of appeal while imposing procedural costs beyond the claimant's means. A building may technically permit alteration while concealing utilities behind surfaces that must be destroyed to locate them.

Definition 5 (Openness gap)**.** The openness gap for agent i is $\Delta_i(x) = B_T(x) - B_i(x) \geq 0$, where B_T is the breadth of $T(x)$.

A large Δ_i indicates that the artifact's advertised revisability is not available to the agent whose freedom supposedly justifies it. Because costs, knowledge, and authority are unevenly distributed, the same artifact may be latently incomplete for one agent and inertly incomplete for another. An undocumented codebase may remain continuable for its original author but inert for every successor. A planning regime may be revisable for landowners with legal counsel but effectively final for tenants affected by its classifications. The unit of analysis is therefore the relation between an artifact and a situated continuer.

2.5 Epistemic, operational, and jurisdictional closure

The situated tuple (9) separates three ways in which a continuation can be withdrawn, and the openness gap decomposes accordingly. Fix an agent i and consider the chain

$$T(x) \;\supseteq\; T_i^{\Sigma}(x) \;\supseteq\; T_i^{\Sigma\mathcal{K}}(x) \;\supseteq\; A_i(x),$$

where T_i^{Σ} retains the formally permitted, identity-preserving states whose paths i can interpret from Σ_t, $T_i^{\Sigma\mathcal{K}}$ further retains those within i's resources under $\mathcal{K}_t$, and A_i further retains those i is able and authorized to select under $\mathcal{W}_t$.

Definition 6 (Three kinds of closure)**.** A continuation $y \in T(x)$ is *epistemically closed* for i if $y \notin T_i^{\Sigma}(x)$: it has ceased to be represented in a form i can interpret. It is *operationally closed* if $y \in T_i^{\Sigma}(x) \setminus T_i^{\Sigma\mathcal{K}}(x)$: it is represented but cannot be executed within i's means. It is *jurisdictionally closed* if $y \in T_i^{\Sigma\mathcal{K}}(x) \setminus A_i(x)$: it is intelligible and affordable, but no authority recognizes i as entitled to select it.

Because breadth is monotone under inclusion, the openness gap telescopes into nonnegative parts,

$$\Delta_i \;=\; \underbrace{(B_T - B_i^{\Sigma})}_{\Delta_i^{\Sigma}} \;+\; \underbrace{(B_i^{\Sigma} - B_i^{\Sigma\mathcal{K}})}_{\Delta_i^{\mathcal{K}}} \;+\; \underbrace{(B_i^{\Sigma\mathcal{K}} - B_i)}_{\Delta_i^{\mathcal{W}}}, \tag{12}$$

and, when the intermediate breadths are positive, the ratio B_i/B_T factors as a product of three ratios, one for each kind of closure. The attribution depends on the order of filtering: a state that is both unintelligible and forbidden is counted here as epistemically closed. The order used follows the sequence in which a continuer typically meets the obstacles, first understanding, then means, then permission. The decomposition explains why transparency alone does not produce revisability. Transparency reduces Δ_i^{Σ} and leaves $\Delta_i^{\mathcal{K}}$ and $\Delta_i^{\mathcal{W}}$ untouched. A deprecated interface can be fully documented and no longer supported. A building's walls can be opened by anyone with a saw while its concealed utilities remain unmapped. A record can be legible and cheap to amend yet amendable only by an office the affected party cannot reach. Continuation rights, the subject of Section 9, are claims on $\mathcal{W}$, and their absence is jurisdictional closure even where epistemic and operational openness are high.

Definition 7 (Public continuability)**.** Let $w_i \geq 0$ be normative weights over a population $I = \{1, \dots, n\}$, Ineq an inequality index on $\mathbb{R}^n_{\geq 0}$ (for example a Gini coefficient scaled by the mean), and $\xi \geq 0$ a concentration penalty. Then

$$\Gamma_{\text{pub}}(x) \;=\; \sum_{i \in I} w_i \Gamma_i(x) \;-\; \xi \; \mathrm{Ineq}\,(\Gamma_1(x), \dots, \Gamma_n(x)). \tag{13}$$

The inequality term matters because a system can possess abundant continuation paths while reserving all of them for a proprietor. Proprietary extension interfaces, discretionary waivers, closed standards, and prohibitively expensive appeals preserve revisability at the system's center while presenting finality at its edges. Such a system is not simply closed. It is *asymmetrically unfinished*: those who govern it may treat its current form as provisional, while those governed by it must treat the same form as complete.

Proposition 3. *An increase in aggregate technical extensibility can coincide with a decrease in public continuability.*

Proof. Let a transformation $x \mapsto x'$ add many private transformation paths for a controlling agent j, so that $|T(x')| > |T(x)|$ and Γ_j increases by $a > 0$, while each noncontrolling agent $i \neq j$ loses $b_i > 0$ (for example because the new paths are coupled to the removal of a public interface). Let $\Delta\,\mathrm{Ineq} \geq 0$ be the induced change in concentration. Then

$$\Gamma_{\text{pub}}(x') - \Gamma_{\text{pub}}(x) \;=\; w_j a - \sum_{i \neq j} w_i b_i - \xi\,\Delta\,\mathrm{Ineq},$$

which is negative whenever $w_j a < \sum_{i \neq j} w_i b_i + \xi \Delta\,\mathrm{Ineq}$, although the technically reachable set has grown. $\square$

2.6 The decay of latent incompletion

The difference between inert and latent incompletion is also temporal. Even a well-structured unfinished artifact loses continuability if it is not maintained. Suppose

$$\frac{d\Gamma}{dt} \;=\; M_t - \eta_x - \eta_\Sigma - \eta_W, \tag{14}$$

where M_t is the continuability restored by maintenance, η_x material degradation, η_Σ loss of interpretive information, and η_W attrition in the community capable of continuing the artifact. If $M_t < \eta_x + \eta_\Sigma + \eta_W$ on an interval, Γ declines there, and if the deficit is bounded below by some $\eta^* > 0$ the artifact reaches $\Gamma(x_{t^*}) = \varepsilon$ no later than $t^* = t_0 + (\Gamma(x_{t_0}) - \varepsilon)/\eta^*$. At t^* latent incompletion crosses into inert incompletion. The crossing need not involve conspicuous failure. File formats become unreadable, contributors disperse, oral knowledge disappears, replacement components cease to exist, procedural exceptions accumulate, and the artifact remains visibly present while its future contracts.

Preservation must therefore include more than retention of the current state. It must preserve the conditions under which a future agent can understand and enter that state. Documentation, provenance, replaceable components, exposed assumptions, test suites, appeal procedures, institutional memory, and interoperable formats are not ancillary conveniences. They are *continuation infrastructure*, the terms of M_t that keep the difference between unfinished and abandoned from being decided by decay. Studies of repair make the same point from the side of practice: maintenance is not a residue of production but a standing condition under which technical systems continue to function at all [16].

This account yields an evaluative test for claims that an unfinished artifact is valuable because it remains open. Is its state legible? Are intervention points practically accessible? Does more than one substantively different continuation remain viable? Does transformation preserve the organization that makes the artifact worth continuing? Is this capacity distributed beyond the original author? Is maintenance sufficient to prevent the option set from expiring? These are not independent virtues to be accumulated without limit; they constrain one another. Excessive closure suppresses diversity, excessive plasticity destroys legibility, excessive preservation immobilizes, and insufficient preservation converts alteration into replacement. Latent incompletion occupies a bounded region between premature closure and disintegrative openness. It is not the absence of decision but a designed relation between what has been decided and what remains available for later determination. That relation must now be priced.

3 The Option Value of the Unfinished

The ordinary economics of completion begins with a project whose benefits remain unavailable until a terminal condition is satisfied. Construction in progress does not shelter, an uncompiled program does not execute, and an unresolved application does not confer the right it requests. Delay therefore appears primarily as a cost. Let V_g be the value of the terminal state once realized, c_t the cost of continued work, and $r > 0$ a discount rate. The project is worth completing at time T when

$$e^{-rT} V_g \;>\; \int_0^T e^{-rt} c_t \, dt. \tag{15}$$

Within this representation incompletion has no positive term. It appears only as foregone use, continuing expenditure, and exposure to the possibility that g is never reached, and the rational actor completes as quickly as resources permit.

This model is appropriate when the terminal state is known, when its desirability is stable, when completion can be reversed without significant loss, and when future information would not alter the preferred design. These conditions often fail. The environment at T may differ from the one anticipated at 0. Those who inherit the artifact may value states its producer never considered. Later evidence may reveal an attractive terminal form to be hazardous, exclusionary, or difficult to maintain. And completion is frequently irreversible in the strong sense studied by the literature on increasing returns: once a design accumulates users, complementary assets, and learning effects, the cost of departing from it grows with its adoption, so that an early and possibly accidental choice becomes locked in [2]. Pierson extends the same mechanism to political institutions, where coordination effects, adaptive expectations, and the asymmetric power conferred by early settlements make paths self-reinforcing [27]. Under such conditions waiting does not merely postpone benefit. It may preserve the right to choose after uncertainty has diminished and before increasing returns have hardened the choice.

3.1 Deferral as an option

Let $S = \{s_1, \dots, s_n\}$ be possible future environments with probabilities $p_1, \dots, p_n$, and let $Y = \{y_1, \dots, y_m\}$ be terminal states reachable from the present unfinished state x. Write

$$Z_j(s_k) \;=\; U(y_j, s_k) - C_k(x, y_j), \tag{16}$$

where $U(y_j, s_k)$ is the value of completed state y_j in environment s_k and C_k the cost of reaching it, allowed to depend on the environment. If completion requires selecting y_j now, the best expected value of immediate completion is

$$V_{\text{now}}(x) \;=\; \max_j \sum_k p_k Z_j(s_k). \tag{17}$$

If the artifact can remain latently incomplete until the environment is revealed, deferred selection yields

$$V_{\text{defer}}(x) \;=\; \sum_k p_k \max_j Z_j(s_k) \;-\; H(x), \tag{18}$$

with $H(x)$ the maintenance cost of the unfinished state. The option value of retained incompletion is

$$O(x) \;=\; V_{\text{defer}}(x) - V_{\text{now}}(x). \tag{19}$$

$O(x)$ does not imply that deferral is always preferable. Maintenance may be costly, delay may withhold urgently needed benefits, and paths may disappear. The point is that incompletion contributes a value which conventional project accounting sets to zero before evaluation begins.

Proposition 4. *If maintaining latent incompletion is costless, all present continuation paths remain available after uncertainty is resolved, and completion is irreversible, then $O(x) \geq 0$, with equality if and only if a single continuation is optimal in every environment of positive probability.*

Proof. With $H = 0$, for every fixed j and every k, $\max_\ell Z_\ell(s_k) \geq Z_j(s_k)$. Multiplying by p_k and summing gives $\sum_k p_k \max_\ell Z_\ell(s_k) \geq \sum_k p_k Z_j(s_k)$ for each j, hence $\geq \max_j \sum_k p_k Z_j(s_k)$, so

$O(x) \geq 0$. If some j^* maximizes $Z_\cdot(s_k)$ for every k with $p_k > 0$, both sides equal $\sum_k p_k Z_{j^*}(s_k)$. Conversely, if equality holds, let j^* attain V_{now}; then $\sum_k p_k[\max_\ell Z_\ell(s_k) - Z_{j^*}(s_k)] = 0$ with nonnegative summands, so $Z_{j^*}(s_k) = \max_\ell Z_\ell(s_k)$ whenever $p_k > 0$. □

Proposition 4 is a standard value-of-information inequality, included to fix notation rather than as a new result. Its role in the economics of irreversible choice is well established. Arrow and Fisher and, independently, Henry showed that when a development decision is irreversible and information will arrive, the option of preservation carries a value, often called quasi-option value, that expected-value comparisons omit [1, 14]; Dixit and Pindyck developed the general theory of investment under irreversibility and uncertainty [10]. What the present framework adds is the reinterpretation of the retained option as a property of an unfinished artifact and of the authority that may withdraw it.

The assumptions are deliberately strict in order to isolate the value of preserved choice. Once maintenance cost, foregone use, and expiring opportunities enter, deferral remains rational only while the marginal informational benefit exceeds the marginal cost of waiting. Let $O(t)$ be the expected option value retained by remaining unfinished until t, $H(t)$ cumulative maintenance cost, and $L_{\text{use}}(t)$ cumulative loss from delayed use. If these are differentiable and O is concave, an interior optimal completion time t^* satisfies

$$O'(t^*) \;=\; H'(t^*) + L'_{\text{use}}(t^*). \tag{20}$$

Before t^* another interval of incompletion produces more expected informational value than it costs; after t^* delay consumes more value than it preserves.

The condition clarifies why pressure to complete is rational in one domain and destructive in another. When an unfinished bridge withholds a necessary crossing and its engineering requirements are well understood, L'_{use} is large and O' small, so prompt completion is justified. When a classification system will determine eligibility for decades while the represented population and the consequences of categorization remain poorly understood, O' may be substantial, and premature closure imposes durable costs exceeding the convenience of immediate finalization.

3.2 Postponing a decision versus preserving a decision space

Uncertainty alone does not create option value; the present artifact must preserve the capacity to act on later information. Let $q_j(t) \in [0, 1]$ be the probability that continuation y_j remains admissible at t. Then

$$V_{\text{defer}}(t) \;=\; \sum_k p_k(t) \max_j \left[q_j(t)U(y_j, s_k) - C_{jk}(t)\right] - H(t) - L_{\text{use}}(t), \tag{21}$$

where $p_k(t)$ is the posterior available at t. If information improves while the q_j decay, knowledge arrives after it can be used. Documentation, reversible joints, replaceable modules, interoperable formats, preserved evidence, and living communities of practice all act on $q_j(t)$; they are the expenditures that keep the option from expiring, which is to say they are the maintenance term M_t of (14) viewed from the side of valuation.

The model reveals the difference between postponing a decision and preserving a decision space. A committee may defer a decision while allowing all but one alternative to disappear. A software project may postpone a release while accumulating dependencies that make architectural revision increasingly costly. A government may call a classification provisional while building benefits, obligations, and databases around it until reversal becomes politically implausible, which is precisely Pierson's mechanism operating under a label of provisionality [27]. Let Y_t be the continuation set at t. *Pure postponement* holds the declaration open while permitting $Y_{t+1} \subsetneq Y_t$. *Structured incompletion* preserves a designated subset $Y^* \subseteq Y_t$ such that $Y^* \subseteq Y_{t+\Delta}$ throughout the interval in which later information is expected to matter. An institution can advertise openness merely by withholding a terminal announcement while its material and procedural commitments eliminate the alternatives the announcement would supposedly decide among.

3.3 Completion as a sequence of commitments

Completion is better represented as a sequence of commitments than as a single terminal event. Let $x_0 \to x_1 \to \ldots \to x_T$ be a development path. Each transition eliminates some continuations and enables others. Define the foreclosure at step t by

$$f_t \;=\; B(x_t) - B(x_{t+1}) + N_t, \tag{22}$$

where $N_t \geq 0$ is the breadth of genuinely new continuations created by the transition, so that the gross contraction $B(x_t) - B(x_{t+1}) + N_t$ counts what was removed rather than the net change. The cumulative foreclosure of the path is

$$F_T^{\mathrm{cum}} \;=\; \sum_{t=0}^{T-1} \max(0, f_t). \tag{23}$$

Two paths may reach the same nominal terminal state with very different F_T^{cum} and, more importantly, with foreclosure distributed differently in time. One path may preserve revisability until late, when more information is available; another may fix consequential choices early and devote the remaining period to polishing their consequences. Both appear equally complete at T, but the second has exposed the project to greater error under uncertainty. The following proposition makes this precise using the standard representation of information as a filtration.

Proposition 5 (Delayed irreversible commitment)**.** *Let $(\mathcal{F}_t)$ be a filtration representing the information available at time t, and let $Z_1, \ldots, Z_m$ be integrable payoffs of mutually exclusive designs. Suppose two development paths have equal direct cost and differ only in the time at which the design is irreversibly fixed, $t_1 < t_2$, and that preserving the choice until t_2 eliminates no design. Then the expected value of fixing at t_2 weakly exceeds that of fixing at t_1.*

Proof. The expected value of committing at t with the best available choice is $W_t = \mathbb{E}[\max_j \mathbb{E}[Z_j \mid \mathcal{F}_t]]$. For each j, by the tower property and monotonicity of conditional expectation,

$$\mathbb{E}[Z_j \mid \mathcal{F}_{t_1}] \;=\; \mathbb{E}\big[\mathbb{E}[Z_j \mid \mathcal{F}_{t_2}] \,\big|\, \mathcal{F}_{t_1}\big] \;\leq\; \mathbb{E}\Big[\max_\ell \mathbb{E}[Z_\ell \mid \mathcal{F}_{t_2}] \,\Big|\, \mathcal{F}_{t_1}\Big].$$

Taking the maximum over j on the left and expectations on both sides yields $W_{t_1} \leq W_{t_2}$. Equality holds when the information arriving between t_1 and t_2 almost surely does not change the optimal design. $\square$

Proposition 5 is the dynamic form of the same inequality and a special case of Blackwell's comparison of experiments, under which a more informative signal, here a refinement of the filtration, is weakly more valuable for every decision problem [5].

Proposition 5 supplies a formal basis for *selective* completion. Components whose early fixation enables all later work may need stabilizing; components whose fixation primarily restricts future use should remain open when preserving them is cheap. The problem is not solved by maximizing delay globally but by ordering commitments according to dependency and reversibility. For components $c_1, \ldots, c_n$ let a_i be the enabling value of fixing c_i, f_i its expected foreclosure cost, and r_i its reversal cost, and define the commitment priority

$$\pi_i \;=\; \frac{a_i}{f_i + r_i}, \tag{24}$$

with the convention $\pi_i = +\infty$ when $f_i + r_i = 0$, so that costless commitments are made as soon as they enable anything. Components with large enabling value and small foreclosure and reversal consequences should be fixed earlier; components with small enabling value and large irreversible consequences later. The ratio is a heuristic, not a decision rule. Its terms may resist precise quantification, and ratio rules fail under dependencies, because a valuable bundle of prerequisites cannot be selected component by component. Section 4 replaces it with an exact optimization over the dependency structure. Its purpose here is to show that "finish the foundation before the walls" is not merely a physical sequence. It is a general principle of commitment architecture: stabilize what makes other work possible, and delay what unnecessarily determines the purposes for which that work may later be used.

3.4 Who owns the option

Option value becomes politically significant when those who receive present benefits differ from those who bear future foreclosure costs. Let immediate completion yield benefits b_{P} to the producer and b_U to current users while imposing expected cost c_F on future users through the lost continuation set. A social evaluation compares $b_{\mathsf{P}} + b_U - c_F$ with the value of maintaining incompletion. If the decision rule contains only the producer's utility, or discounts future users at an extreme rate, completion is chosen whenever $b_{\mathsf{P}} > 0$ regardless of c_F. The institution then behaves as though future continuability had no owner.

The option model above hides the same problem inside its maximum. Writing $\max_j Z_j(s_k)$ assumes that whoever exercises the option is informed, competent, and aligned with the valuation being computed, which is exactly what cannot be assumed politically. An option controlled by a monopolist can be valuable to the controller and harmful to everyone it affects. The primitive should therefore index option value by the affected agent and by the controller.

Let $\mathfrak{A} \subseteq \Omega$ be the set of *presently admissible* states: those not excluded by declared safety constraints, binding legal constraints, or other limits that the present valuation is not authorized to override. Catastrophic continuations are removed here, before breadth or value is computed. Admissibility is itself relative to a level of authority. A legal constraint binding on the present valuation may be a revisable closure at another institutional level, so $\mathfrak{A}$ is fixed for the purposes of one decision without being placed beyond revision altogether. Let $u_i(y, s)$ be the value of state y to agent i in environment s, with no sign restriction, and let $A^j(x) \subseteq \mathfrak{A}$ be the admissible continuations that agent j is able and authorized to select from x. When j selects according to its own valuation after observing the environment, write $y^j(x, s) \in \arg\max_{y \in A^j(x)} u_j(y, s)$.

Definition 8 (Control-indexed option value)**.** The value to agent i of the option on x controlled by agent j is

$$O_i^j(x) \;=\; \mathbb{E}_s\big[u_i(y^j(x,s), s)\big] \;-\; \max_{y \in A^j(x)} \mathbb{E}_s\big[u_i(y, s)\big], \tag{25}$$

the difference between what i receives when j chooses after learning the environment and what i would receive if the continuation best for i in expectation were fixed now.

Proposition 6 (Whose option)**.** *For every agent j, $O_j^j(x) \geq 0$. For $i \neq j$, $O_i^j(x)$ is nonnegative whenever u_i and u_j induce the same ordering of $A^j(x)$ in every environment, and it can be negative otherwise.*

Proof. $O_j^j \geq 0$ is Proposition 4 applied to the payoffs $u_j(y, \cdot)$, $y \in A^j(x)$, with zero maintenance cost. If u_i and u_j are ordinally aligned on $A^j(x)$ in every environment, then $y^j(x,s)$ also maximizes $u_i(\cdot, s)$, so $O_i^j = O_i^i \geq 0$. For negativity, take two admissible states with $u_i(y_1, \cdot) = 1$ and $u_i(y_2, \cdot) = 0$ in every environment, while $u_j(y_2, s) > u_j(y_1, s)$ in every environment. Then j always selects y_2, and $O_i^j = 0 - 1 = -1$. □

Separating descriptive breadth from valuation removes a sign ambiguity that would otherwise run through the essay. Breadth is unsigned and says nothing about desirability. Option value is signed, is computed only over admissible states, and depends on who controls the choice. A foreclosure that shrinks $A^j(x)$ to $A^j(x')$ costs agent i the amount $\mathbb{E}_s[u_i(y^j(x,s),s)] - \mathbb{E}_s[u_i(y^j(x',s),s)]$, which is nonnegative when $j = i$ and may be negative when the controller is misaligned with i: removing a branch from a misaligned controller's option can benefit those it affects. Several later results are statements about the distribution of O_i^j across the index pair. Proposition 3 describes gains to a controlling agent accompanied by losses to others, and Proposition 14 in Section 9 describes the same pattern for authority over institutions. The producer-side accounting boundary of Observation 1 is the restriction of a social evaluation $\sum_i w_i O_i^j$ to agents i and controllers j inside the producing authority.

The asymmetry is common because present completion is visible while future foreclosure is counterfactual. The completed building can be photographed, the standardized database generates reports, the resolved case can be counted. The alternative room arrangement, the unrecorded category, the later claimant excluded by the finished system leave no corresponding object. Their absence does not appear as a loss because it never becomes actual enough to enter the ledger. Let $M(z) \in [0,1]$ be the probability that consequence z enters institutional evaluation, with realized outputs Z_{real} and lost continuations Z_{lost}. The measured value of completion is

$$\widehat{V}(C) \;=\; \sum_{z \in Z_{\text{real}}} M(z)v(z) - \sum_{w \in Z_{\text{lost}}} M(w)v(w), \tag{26}$$

whereas the true value sets $M \equiv 1$. When $M \approx 1$ on realized outputs and $M \approx 0$ on lost continuations, completion acquires an evidentiary advantage: it produces countable achievements, while foreclosure produces structured nonoccurrences.

Observation 2. *If realized gains are measured more reliably than foreclosed alternatives, an institution acting on measured net benefit will systematically overselect completion.*

Proof. Let true gain be $G > 0$ and true foreclosure loss $L > 0$, measured as $\widehat{G} = m_G G$ and $\widehat{L} = m_L L$ with $1 \geq m_G > m_L \geq 0$. The institution completes when $m_G G > m_L L$, that is when $G/L > m_L/m_G$. Completion is truly beneficial only when $G/L > 1$. Since $m_L/m_G < 1$, every project with $G/L \in (m_L/m_G, 1)$ is completed at a true loss, and this interval is nonempty. Its width $1 - m_L/m_G$ grows as the measurability gap widens. □

Retained openness therefore requires representation. If an alternative cannot be described, its loss cannot be evaluated; if a future claimant cannot be represented in present procedure, the narrowing of that claimant's options appears costless. Maintaining latent incompletion requires a *ledger of possibilities* beside the ledger of achievements. Such a ledger need not enumerate every future use, which would reproduce the fantasy of complete anticipation. It must record where commitments are made, which branches they remove, who is authorized to make them, what evidence would justify reopening them, and what infrastructure keeps reopening possible. Its function in the terms of Observation 2 is to raise m_L.

3.5 An incompletion diagnostic

The unfinished artifact is not valuable because delay is virtuous. Its value arises from the conjunction of uncertainty, irreversibility, and preserved choice. When uncertainty is negligible, commitment loses little. When reversal is cheap, premature completion can be corrected. When no meaningful alternatives remain, unfinishedness is mere delay. Let σ, R, and $D_{\mathfrak{A}}$ be relevant uncertainty, irreversibility, and the diversity of admissible continuations (states in $\mathfrak{A}$ only), and let H and L_{use} be the costs of maintaining openness and of delayed use over the horizon in question, all normalized to $[0, 1]$ so that the ratio is dimensionless. Define

$$\Lambda(x) \;=\; \frac{\sigma\, R\, D_{\mathfrak{A}}}{H + L_{\text{use}} + \epsilon}, \qquad \epsilon > 0. \tag{27}$$

The delayed-use term is essential. Without it, a medically necessary but highly irreversible intervention would score as a strong candidate for remaining unfinished even when delay is itself catastrophic. Λ is an ordinal heuristic, not a measurement. Large Λ indicates a strong prima facie case for preserving latent incompletion; small Λ indicates that closure is unlikely to destroy valuable optionality. The multiplicative numerator is not arbitrary. Proposition 4 shows that option value vanishes when one design dominates in every environment (no effective uncertainty, $\sigma = 0$) or when only one admissible design exists ($D_{\mathfrak{A}} = 0$), and the option has no value to protect when an error can be undone at no cost ($R = 0$), since one may then complete now and switch later. The expression is diagnostic rather than mechanical. Its contribution is to show why uncertainty without irreversibility, irreversibility without alternatives, and alternatives without affordable preservation do not independently justify delay.

Completion should consequently be understood as the exercise of an option, not the natural disappearance of deficiency. Once exercised, the option may cease to exist. The normative issue is not whether decisions must be made, since action always selects among futures, but whether each decision occurs at a time, scale, and level of abstraction appropriate to the evidence available. The next problem is architectural in the broadest sense: how an artifact can stabilize enough structure to function while withholding precisely those commitments that later agents have the strongest claim to determine for themselves.

4 Selective Completion and the Architecture of Commitment

The preceding sections established that closure has costs as well as benefits and that its timing matters. They did not say *what* to close. The constructive principle of this essay is that a system should be completed at the levels that enable dependable participation and left unfinished at the levels that determine revisable purposes. This section gives that principle a precise form.

The key observation is that commitments are not independent. Fixing a component is only meaningful if the components it relies upon are themselves fixed, and the value of fixing a component consists largely in the dependable work it makes possible elsewhere. The selection of closures is therefore a problem over a dependency structure, not a list of separate decisions.

4.1 Dependency graphs and effective closure

Model an artifact as a finite directed graph $G = (V, \mathcal{E})$ whose vertices are components or commitments and whose edges $i \to j$ indicate that j depends on i: the behavior or meaning of j is determinate only relative to a determinate i. Write $\mathrm{anc}(j)$ for the set of strict ancestors of j and $\mathrm{desc}(j)$ for its strict descendants. A foundation, a load path, a file format, a unit of measure, and a legal definition sit near the sources of such a graph; a room's function, an application's workflow, a policy's substantive target, and a curriculum's choice of texts sit near its sinks.

Many of the objects of this essay have circular dependencies. Legal categories shape administrative evidence, which reshapes the categories; interfaces constrain modules, while accumulated module practice changes the interface; recommendations shape the preferences on which they are trained; purposes determine which foundations are needed, while existing foundations determine which purposes appear possible. The analysis therefore does not assume that the underlying system is acyclic. It works on the *condensation* of G, obtained by collapsing each strongly connected component into a single vertex whose weight is the sum of its members' weights. The condensation is always acyclic. Each of its vertices is a bundle of mutually dependent commitments that can only be fixed or left open together, and when every strongly connected component is a singleton the condensation is G itself. Throughout this section G denotes the condensation, and every result below applies to it unchanged. Closure is then decided over bundles, which is the level at which it is in fact decided when commitments co-constitute one another.

A closure operator fixes a subset $S \subseteq V$ and leaves $V \setminus S$ revisable. Not every subset corresponds to an effective closure. If j is declared fixed while some $i \in \mathrm{anc}(j)$ remains fluid, then the behavior of j still varies with i, and the declaration is nominal rather than effective: users who rely on j inherit the variability of i without the right to revise either. This motivates the following requirement.

Definition 9 (Effective closure)**.** A set $S \subseteq V$ is an *effective closure* if it is an order ideal of G: whenever $j \in S$ and $i \to j$, then $i \in S$. Equivalently, $\mathrm{anc}(j) \subseteq S$ for every $j \in S$. Denote the family of effective closures by $\mathcal{I}(G)$.

$\mathcal{I}(G)$ is closed under union and intersection and therefore forms a distributive lattice with least element $\emptyset$ (nothing fixed) and greatest element V (everything fixed). The two extremes correspond to the positions this essay rejects: disintegrative openness and total completion.

4.2 Weights

Assign to each $j \in V$ three quantities. The *enabling value* $E_j \geq 0$ is the value of dependable work made possible by fixing j. A natural specification apportions participation value across dependencies: if each vertex v supports work of value ω_v that can be performed dependably once its ancestors are fixed, set

$$E_j \;=\; \sum_{v\,:\,j \in \mathrm{anc}(v) \cup \{v\}} \frac{\omega_v}{|\,\mathrm{anc}(v)| + 1}, \tag{28}$$

so that each unit of participation value is shared equally among the commitments on which it rests. Vertices with many and valuable descendants receive high enabling value; call this their *enabling centrality*. A sink receives only its share of its own direct use. The *foreclosure cost* F_j is the value, to the affected population, of the identified admissible alternatives that closing j removes, assessed with the information available at closure and computed with the control-indexed valuation of Definition 8. By Proposition 6 it is nonnegative when those who would exercise the alternatives are aligned with the affected population, and it may be negative when they are not. The *reversal cost* $R_j \geq 0$ is the cost of undoing the closure, and the *expected reversal burden* is $\bar{R}_j = \Pr(\text{closure of } j \text{ proves mistaken}) \cdot R_j$. Keeping F_j narrow in this way separates the loss of known alternatives from the risk of error under unresolved uncertainty, which $\bar{R}_j$ carries, so that the two are not counted twice. The net weight of closing j is

$$w_j \;=\; E_j - F_j - \bar{R}_j, \tag{29}$$

and the value of an effective closure is

$$J(S) \;=\; \sum_{j \in S} w_j, \qquad S \in \mathcal{J}(G). \tag{30}$$

The linear objective is an approximation, since interactions among commitments can make foreclosure costs nonadditive, but it already captures the essential structure, and Section 12 extends the weights to incorporate uncertainty and distribution.

4.3 Optimal closure is a maximum-weight ideal

Proposition 7. *The problem* $\max_{S \in \mathcal{J}(G)} J(S)$ *has a solution. The set of optimal effective closures is closed under union and intersection, so there exist a unique minimal optimal closure* S_* *and a unique maximal optimal closure* S^**. The problem is solvable in time polynomial in* $|V|$*.*

Proof. $\mathcal{J}(G)$ is finite and nonempty, so a maximum exists. J is modular: for any S, T, $J(S \cup T) + J(S \cap T) = J(S) + J(T)$, because each vertex is counted the same number of times on both sides. If S and T are optimal with value $J^\star$, then $S \cup T$ and $S \cap T$ are ideals, each with value at most $J^\star$, and their sum is $2J^\star$, so both equal $J^\star$. Hence optimal closures form a sublattice, whose meet and join are S_* and S^*. Finally, maximizing a vertex-weighted sum over the ideals of a directed graph is the classical maximum-weight closure problem (with edges reversed), which reduces to a minimum s–t cut on a network with $|V| + 2$ nodes and is therefore polynomial [26]. □

The lattice structure has an interpretive consequence. S_* is the most open arrangement consistent with optimal value; S^* the most closed. Where the two differ, closure is a matter of indifference in aggregate value, and Section 12 will argue that such ties should be broken toward S_* whenever the vertices in $S^* \setminus S_*$ determine purposes rather than foundations.

Proposition 8 (Local optimality conditions)**.** *Let* S *be an optimal effective closure. Then*

(i) every descendant-closed *subset* $U \subseteq S$ *(that is,* $j \in U$ *and* $j \to k$ *with* $k \in S$ *imply* $k \in U$*) satisfies* $\sum_{j \in U} w_j \geq 0$*;*

(ii) every ancestor-closed extension $U \subseteq V \setminus S$ *(that is,* $S \cup U \in \mathcal{J}(G)$*) satisfies* $\sum_{j \in U} w_j \leq 0$*.*

Proof. For (i), $S \setminus U$ is again an ideal: if $k \in S \setminus U$ and $i \to k$, then $i \in S$, and $i \notin U$ since otherwise descendant-closure would force $k \in U$. Optimality gives $J(S \setminus U) \leq J(S)$, i.e. $\sum_U w_j \geq 0$. For (ii), $S \cup U$ is an ideal by assumption, and optimality gives $J(S \cup U) \leq J(S)$. □

Two corollaries formalize the difference between foundations and purposes.

Corollary 1 (Purposes remain open)**.** *Call a sink of G a* purpose vertex*. If a purpose vertex j has $w_j < 0$, that is, if its enabling value falls short of its foreclosure and reversal costs, then j belongs to no optimal closure; if $w_j = 0$, $j \notin S_*$. In particular, if any purpose vertex has negative net weight, total completion $S = V$ is suboptimal.*

Proof. A sink j in an ideal S forms a descendant-closed singleton, so by Proposition 8(i) $w_j \geq 0$ whenever j lies in an optimal S. If $w_j = 0$, removing j preserves optimality, so the minimal optimum excludes it. □

Corollary 2 (Foundations are closed at a price)**.** *A vertex i with $w_i < 0$ may belong to an optimal closure when it is a prerequisite of a descendant set whose aggregate contribution compensates for w_i; it belongs to every optimal closure exactly when $i \in S_*$. Conversely, if some source vertex has $w_i > 0$, then the empty closure is suboptimal.*

Proof. Proposition 8(i) constrains the total weight of a descendant-closed set within an optimum, not the weights of its members, so negative members are permitted when their descendants compensate. Since the optimal closures form a lattice with least element S_* (Proposition 7), i lies in every optimum if and only if $i \in S_*$. For the second, $\{i\}$ is an ideal when i is a source, and $J(\{i\}) = w_i > 0 = J(\emptyset)$. □

Corollary 2 explains a familiar pattern: foundations are often expensive to fix, foreclose options, and are costly to reverse, and are fixed anyway because the dependable work built on them repays the price. What the corollary adds is the condition under which this is justified. A foundation is warranted by the positive aggregate weight of what it supports, not by its position in the construction sequence. A foundation that supports only purposes the system has no standing to settle does not inherit their justification.

The identification of sources with foundations and sinks with purposes is a useful approximation, not a definition. A purpose can sit upstream when it determines the architecture of everything below it, as a mission determines an organization's structure, and a component can sit downstream while serving as infrastructure for later users, as a data format does once other systems consume it. Foundation and purpose are better defined by role in the present decision. A component is foundation-like to the extent that its enabling centrality is high relative to its foreclosure and reversal burden, and purpose-like to the extent that uncertainty about its use is high, its closure is sensitive to who holds authority, and its enabling value is small. Graph position is evidence for these judgments, not their definition. Corollary 1 is stated for sinks because a sink has no dependents to compensate for its cost; by Proposition 8(i) the same argument applies to any descendant-closed set of purpose-like components whose total weight is negative. Conversely, a structurally upstream component on which nothing of value depends is not a foundation in the relevant sense, and it inherits no justification from its position.

The resulting principle is that neither maximal closure nor maximal openness generally maximizes value. Under the specification (28), sources tend to have high enabling centrality while sinks have little, and if foreclosure is concentrated where uncertainty about future use is greatest, which is typically at the level of purposes, the optimum is an interior ideal: firm near the sources, open near the sinks. Baldwin and Clark describe the same partition in the design of modular systems as the separation between *visible design rules*, fixed early and binding on every module, and *hidden design parameters*, left to the discretion of whoever works within a module [3]. The present result states when that partition is optimal rather than merely conventional, and it places the boundary by the weights rather than by habit.

4.4 Closure over time

Combining Proposition 7 with Proposition 5 yields a temporal picture. Let the weights $w_j(t)$ evolve as information arrives. A development process that closes $S_*(t)$ at each t produces a *closure schedule*. Whether the schedule is monotone, so that nothing is fixed only to be unfixed at reversal cost, depends on how the weights move.

Lemma 1 (Nested closure)**.** *If $w_j(t)$ is nondecreasing in t for every j, then the minimal optimal closures are nested: $t < t'$ implies $S_*(t) \subseteq S_*(t')$.*

Proof. Let $S = S_*(t)$, $S' = S_*(t')$, and write J_t for the objective with weights $w(t)$. By modularity,

$$J_t(S) - J_t(S \cap S') = \sum_{j \in S \setminus S'} w_j(t) \ \le \ \sum_{j \in S \setminus S'} w_j(t') = J_{t'}(S \cup S') - J_{t'}(S') \ \le \ 0,$$

the last inequality by optimality of S' at t'. Hence the ideal $S \cap S'$ is optimal at t, and minimality of S gives $S \subseteq S \cap S' \subseteq S'$. □

The hypothesis is substantive. Foreclosure costs typically fall as uncertainty resolves, which raises weights, but enabling values can also fall when needs shift, and pointwise reoptimization would then prescribe unfixing. Where weights are not monotone, nesting is a design policy rather than a consequence of optimization: one commits to a vertex only when its weight is expected to remain positive in the context of its dependents. The priority ratio (24) is a heuristic for traversing such a schedule. The ideal structure is what makes it coherent, because it forbids fixing a purpose before the foundations that give it meaning and forbids treating a foundation's early fixation as evidence that its dependents must also be fixed.

5 Human Development and the Temporality of Plasticity

The same structure appears in developing agents, where it has been argued for independently. Bruner proposed that the long immaturity of human young is not a delay before competence but an adaptation: extended dependence creates a protected period in which behavior can be assembled, disassembled, and recombined in play at low cost, so that the eventual repertoire is broader than one fixed early could be [8]. Karmiloff-Smith argued that the modular organization of adult cognition is in large part a *product* of development rather than its starting point: domain-general processes become progressively specialized, and representations are repeatedly redescribed into more explicit and more flexible formats [20]. In the vocabulary of Section 4, development is a closure schedule. Some structure is fixed early to make learning possible at all; much is fixed late, after experience has reduced uncertainty about what the environment requires.

Instantiation. The state Ω is the agent's repertoire of competences and dispositions; the recognized continuations A are the competences still learnable from it; the information Σ is the record of behavior together with the conditions under which it was produced; the cost structure $\mathcal{K}$ is the effort and time of learning; the authority $\mathcal{W}$ is divided among the agent, caregivers, and institutions that assign environments; and $\sim_{\mathcal{R}}$ is the continuity of the person across development. The models below are stylized. They isolate mechanisms and do not explain development empirically.

5.1 The plasticity–stability trade-off

Let $\mathcal{P} \in [0,1]$ be the fraction of an agent's adaptive capacity held plastic, available to incorporate new environmental information, and $\mathcal{S} = 1-\mathcal{P}$ the fraction held stable, retaining what has been learned. Adaptation requires both. Define developmental viability

$$V(\mathcal{P}) \;=\; \mathcal{P}^{\alpha}(1-\mathcal{P})^{\beta}, \qquad \alpha, \beta > 0. \tag{31}$$

Observation 3. *V attains its unique maximum at $\mathcal{P}^* = \alpha/(\alpha+\beta) \in (0,1)$. In particular neither maximal plasticity nor maximal fixity is optimal.*

Proof. $\log V = \alpha \log \mathcal{P} + \beta \log(1-\mathcal{P})$ is strictly concave on $(0,1)$ and tends to $-\infty$ at both endpoints. Its derivative $\alpha/\mathcal{P} - \beta/(1-\mathcal{P})$ vanishes only at $\mathcal{P}^*$. □

The observation is a property of the chosen specification, not an empirical finding: the interior optimum has been placed in the functional form, and the formalism contributes only an exact statement of the trade-off it encodes. It is the developmental counterpart of Proposition 2. Openness to formation is not passive indeterminacy: an agent with $\mathcal{P} = 1$ retains nothing and therefore cannot build on what it encounters, while an agent with $\mathcal{P} = 0$ cannot encounter anything new.

The exponents need not be constant. The value of plasticity depends on how much remains to be learned about the environment. Let σ_t^2 be the agent's residual uncertainty about environmental requirements at time t, nonincreasing as experience accumulates, and suppose the plasticity exponent is $\alpha_t = \alpha_0\, g(\sigma_t^2)$ for some increasing g. Then

$$\mathcal{P}^*(t) \;=\; \frac{\alpha_0 g(\sigma_t^2)}{\alpha_0 g(\sigma_t^2) + \beta} \tag{32}$$

is nonincreasing in t. The optimal developmental schedule holds capacity open while uncertainty is high and closes it progressively as uncertainty resolves. This is Proposition 5 applied to a learner: irreversible specialization should follow, not precede, the information that would justify it. The model suggests one reading of critical periods, as points at which the marginal informational value of continued plasticity falls below its cost in unconsolidated competence. This is an interpretation offered by the model, not an explanation established by it. Critical periods also reflect maturational timing, metabolic constraint, circuit stabilization, and developmental cascades that the model does not represent.

5.2 Premature specification

Tools, schools, and institutions frequently infer an agent's settled capacities or preferences from behavior produced during an unfinished process and then fix an environment accordingly: a track, a stream, a diagnosis, a recommended career. Such inferences have two costs. The first is ordinary estimation error, which is larger early. The second is subtler: once the environment is fixed, subsequent behavior is produced partly by the environment, and evidence that the inference was wrong becomes harder to observe.

Consider an agent of latent type $\theta \in \{a, b\}$ and an institution that assigns an environment $k \in \{a, b\}$ at time t_0 on the basis of observations $y_1, \dots, y_{t_0}$ with $y_s = \mu_\theta + \xi_s$, $\mu_a = \Delta_\theta/2$, $\mu_b = -\Delta_\theta/2$, $\Delta_\theta > 0$, and independent Gaussian noise of variance σ_ξ^2. Assignment by the sign of the sample mean misassigns with probability

$$\epsilon(t_0) \;=\; \Phi\left(-\frac{\Delta_\theta \sqrt{t_0}}{2\sigma_\xi}\right), \tag{33}$$

where Φ is the standard normal distribution function, which is decreasing in t_0. Early assignment is therefore more error-prone, and if misassignment reduces the agent's learnable breadth by $\Delta B > 0$ the expected loss is $\epsilon(t_0)\Delta B$, again decreasing in t_0. This part is familiar. Now let post-assignment behavior depend on the environment: after assignment, $y_s = \mu_\theta + \gamma(\mathbf{1}[k = a] - \mathbf{1}[k = b]) + \xi_s$, where $\gamma \geq 0$ is the effect of the environment on the observed measure (an agent placed in a lower track receives less of the instruction the measure rewards).

Proposition 9 (Self-confirming specification)**.** *Suppose the institution evaluates post-assignment behavior with the pre-assignment likelihood, ignoring the environmental effect γ, and consider a misassigned agent.*

(i) If $\gamma > \Delta_\theta/2$, the agent's expected post-assignment signal is closer to the mean of the assigned (wrong) type than to that of the true type, so the expected evidence favors the assignment.

(ii) If $\gamma \geq \Delta_\theta$, the expected signal is at least as extreme as the pre-assignment mean of the assigned type, so the misassigned agent is on average classified into that type at least as strongly as a correctly assigned agent observed before assignment.

Proof. Take $\theta = a$ and $k = b$. The expected post-assignment signal is $\Delta_\theta/2 - \gamma$. It is closer to $\mu_b = -\Delta_\theta/2$ than to $\mu_a = \Delta_\theta/2$ exactly when it is negative, that is when $\gamma > \Delta_\theta/2$, which gives (i). It is at most μ_b exactly when $\gamma \geq \Delta_\theta$, and since the Gaussian log-likelihood ratio in favor of b is decreasing in the signal, this gives (ii). The case $\theta = b$, $k = a$ is symmetric. □

Part (i) says that the evidence points the wrong way; part (ii) says that the misassigned agent becomes, on average, indistinguishable from the type to which it was assigned. The proposition identifies the mechanism by which premature closure of a developing agent becomes invisible: the closure produces the behavior that is then cited as its justification. The same loop reappears in Section 7 for predictive systems, where it operates at population scale. The remedy suggested by the formalism is not to abolish assignment but to preserve legibility of the counterfactual. An institution that records the environmental effect γ, or that keeps periodic evaluation environments common to all tracks, keeps the assignment revisable in the sense of Section 2: it maintains the information Σ_t without which a later continuer cannot tell whether the specification was warranted.

Development thus illustrates the general thesis with unusual clarity. A learner is finished well when stable competences support further learning, and finished badly when specification outruns evidence and then manufactures the evidence it lacked.

6 Tools, Interfaces, and Attachment Points

The term *attachment point* has so far been used informally for a place at which a later participant can enter an artifact. This section gives it a definition, distinguishes nominal from effective attachment points, and states conditions under which modularity increases continuability.

Instantiation. Here Ω is the configuration space of an artifact with interfaces; A is the set of transitions producible through them; Σ is documentation and specification; $\mathcal{K}$ includes understanding, implementation, switching, and replacement costs; $\mathcal{W}$ distinguishes the designer's discretion from permissions that hold without it; and $\sim_{\mathcal{R}}$ is the artifact's identity across substitutions of components.

6.1 Definitions

Definition 10 (Attachment point)**.** An attachment point of an artifact is a pair $(\iota, \mathcal{T}_\iota)$, where ι is an interface, a designated boundary in the dependency graph across which a component may be added, substituted, or reconfigured, and $\mathcal{T}_\iota$ is the set of state transitions an external agent can produce through ι *without obtaining discretionary authorization from the original designer.*

The last clause is what distinguishes an attachment point from a request. A form that asks the manufacturer to enable a feature is not an attachment point; a documented socket that accepts any conforming part is. The nominal set $\mathcal{T}_\iota^{\text{nom}}$ contains the transitions permitted by the interface specification and its license. The effective set $\mathcal{T}_\iota^i$ for agent i contains those transitions that also satisfy the conditions of (10): they are documented well enough to be interpreted within tolerance τ, interoperable with components i can obtain, affordable within κ once switching costs are included, accessible under the rights i actually holds, and identity preserving under $\sim_{\mathcal{R}}$.

Definition 11 (Accessibility coefficient)**.** The accessibility coefficient of attachment point ι for agent i is

$$\chi_i(\iota) \;=\; \frac{B(\mathcal{T}_\iota^i)}{B(\mathcal{T}_\iota^{\text{nom}})} \;\in [0,1], \tag{34}$$

the proportion of theoretically permitted extension breadth that is practically available to i, with the convention $\chi_i(\iota) = 1$ when the nominal breadth is zero, since there is then nothing permitted for i to miss.

The coefficient inherits the normalization of B. Using the decomposition $B = \mathsf{m}^2\bar{\delta}$ of Section 1, it can be split into two more informative quantities: the fraction of nominal continuation mass that is reachable by i, and the ratio of the mean differentiation of the reachable subset to that of the nominal set. The first says how much of the permitted space i can reach; the second says whether what i can reach is representative of its diversity.

The coefficient localizes the openness gap of Section 2 to a single interface. Repairable machinery with published schematics, standard fasteners, and available parts has χ close to one for a broad population; the same machine with glued enclosures, serialized components, and unpublished service manuals may be nominally repairable, since nothing forbids opening it, while χ approaches zero for everyone outside an authorized network.

6.2 When modularity helps

Modularity is often assumed to increase openness. It does so only under a condition on understanding costs. Let an artifact of size n be either monolithic or decomposed into m modules separated by interfaces. Consider a potential change Δ with value $v(\Delta)$ to agent i and implementation cost $k(\Delta)$, and let R be the cost of replacing the artifact outright. Let u be the cost to i of understanding the system well enough to make the change safely. The change is *continuable* through the artifact when

$$u + k(\Delta) \;\leq\; \min\{v(\Delta),\, R\}, \tag{35}$$

since otherwise either the change is not worth making or replacement is cheaper than continuation, in which case the artifact supplies material rather than a future. Write $\mathcal{C}(u)$ for the set of continuable changes. Clearly $\mathcal{C}$ is antitone: $u \leq u'$ implies $\mathcal{C}(u') \subseteq \mathcal{C}(u)$.

For a monolith, $u = \Upsilon(n)$ with Υ increasing. For a modular system in which changes are local to one module, $u = \Upsilon(n/m) + c(m)$, where $c(m)$ is the cost of understanding the interfaces and design rules, increasing in m.

Proposition 10 (Modularity and continuability)**.** *Suppose the same local changes are available under both architectures, so that modularization alters their costs but not the set of possible modifications. Then modularization weakly enlarges the set of continuable local changes if and only if $\Upsilon(n/m)+c(m) \leq \Upsilon(n)$, and it enlarges it strictly exactly when some change has $\Upsilon(n/m)+c(m)+k(\Delta) \leq \min\{v(\Delta), R\} < \Upsilon(n)+k(\Delta)$. If $\Upsilon(n/m)+c(m) > R - \min_{\Delta} k(\Delta)$, the modular system admits no continuable changes at all: every modification is cheaper as replacement.*

Proof. The first two claims follow from the antitonicity of $\mathcal{C}$ and the definition of (35). For the third, every Δ then has $u + k(\Delta) > R \geq \min\{v(\Delta), R\}$. □

In real systems modularization also changes which modifications exist, adding some and forbidding others; the proposition isolates the cost channel. If Υ is convex with $\Upsilon(0) = 0$ and c is increasing and convex, the understanding cost $\Upsilon(n/m) + c(m)$ is minimized at an interior m^*: too few modules leave each one opaque, and too many make the design rules themselves the obstacle. Modularity, like malleability in Proposition 2, has an interior optimum.

A second advantage of modularity concerns option value directly. Baldwin and Clark observe that a modular design converts one option on the system as a whole into a portfolio of options on its modules, and that the portfolio is worth more [3]. The underlying inequality, that a portfolio of options is worth at least as much as an option on the portfolio, is a classical result of option pricing [23]; its short proof is included to show how it applies to attachment points.

Proposition 11 (Portfolio of options)**.** *Let $X_1, \dots, X_m$ be integrable random improvements available by redesigning modules $1, \dots, m$, each adoptable only if positive. The value of independently adoptable module options weakly exceeds the value of a single option to adopt all redesigns together:*

$$\sum_{k=1}^{m} \mathbb{E}[\max(X_k, 0)] \;\geq\; \mathbb{E}\Big[\max\Big(\sum_{k=1}^{m} X_k,\, 0\Big)\Big],$$

with strict inequality whenever, with positive probability, the X_k are not all of the same sign.

Proof. Pointwise, $\max(\sum_k X_k, 0) \leq \sum_k \max(X_k, 0)$, since the right side is nonnegative and at least $\sum_k X_k$. Taking expectations gives the inequality. If on an event of positive probability some $X_k > 0$ and some $X_\ell < 0$, then on that event $\sum_k \max(X_k, 0) > \max(\sum_k X_k, 0)$, and the inequality is strict. □

Proposition 11 is the precise sense in which modular attachment points preserve option value: they let later agents exercise the good options without being forced to take the bad ones bundled with them. Proposition 10 states the price. The portfolio exists only for agents whose understanding costs are low enough to reach it.

6.3 Boundary objects as distributed attachment

Star and Griesemer introduced the *boundary object* to describe artifacts, such as specimens, maps, and standardized forms, that are shared among communities with different aims: plastic enough to be adapted to each community's local needs, robust enough to maintain a common identity across them [31]. In the present terms a boundary object is an artifact whose equivalence relation $\sim_{\mathcal{R}}$ is fixed and shared, so that all communities agree it is the same object, while each community holds an attachment point with high accessibility coefficient through which it continues the object in its own direction. It is asymmetric completion distributed across social worlds: closed exactly where closure is needed for coordination and open where each participant must be free to use the object for purposes the others do not share.

The same structure recurs across materially different domains. Habraken's supports fix structure, services, and access while leaving infill to occupants [13]. Brand's slow layers of site and structure carry fast layers of services, space plan, and stuff that change without disturbing what lies beneath [7]. An application programming interface fixes signatures and semantics while leaving implementations and uses open. An open standard fixes a representation while permitting competing implementations. A reproducible scientific method fixes procedures and materials precisely enough that others can extend, vary, or refute the result: the methods section is an attachment point, and a paper without one is inertly complete however polished its conclusions. In each case what is finished is the interface, and what is left unfinished is everything the interface makes possible.

7 Helpful Completion and the Capture of Intention

The artifacts considered so far are completed by their producers. A newer class of artifacts completes *its users.* Autocomplete finishes sentences, recommendation systems finish choices, route planners finish journeys, predictive interfaces finish forms, and generative assistants finish drafts, arguments, and images. These systems act on an object that was until recently treated as the paradigm of latent incompletion: an intention that has not yet settled. The question of this section is when such completion assists an agent and when it replaces the agent's own continuation with the system's.

Instantiation. The state Ω is the agent's space of possible actions together with the purposes they serve; the continuation relation A is the distribution of actions still open to the agent; Σ is what the system and the agent can observe of unassisted behavior; $\mathcal{K}$ includes the refusal cost defined below; $\mathcal{W}$ divides authority over purpose between agent and system; and $\sim_{\mathcal{R}}$ is the identity of the agent's purpose across alternative executions.

7.1 Suggestion-induced divergence

Let $\mathcal{A}$ be a finite set of possible actions and p_0 the distribution over $\mathcal{A}$ describing what an agent would do unassisted. A completion system offers a suggestion $s \in \mathcal{A}$, after which the agent acts according to $p(\cdot \mid s)$. The strength with which the suggestion reorganizes the action space is

$$I_s \;=\; D_{\mathrm{KL}}(p(\cdot \mid s) \,\|\, p_0) \;=\; \sum_{a \in \mathcal{A}} p(a \mid s) \log \frac{p(a \mid s)}{p_0(a)}. \tag{36}$$

$I_s = 0$ exactly when the suggestion changes nothing. Large I_s is not in itself objectionable, since a correction of a typing error changes behavior a great deal. What matters is *where* the divergence falls.

Decompose each action as $a = (\varpi, e)$, where ϖ is the purpose the agent pursues and e the execution by which it is pursued: which word is meant versus how it is spelled, where one is going versus which streets one takes, what one wants to argue versus how the sentence is phrased. By the chain rule for relative entropy,

$$I_s \;=\; \underbrace{D_{\mathrm{KL}}(p(\varpi \mid s) \,\|\, p_0(\varpi))}_{I_s^{\mathrm{purpose}}} \;+\; \underbrace{\mathbb{E}_{\varpi \sim p(\cdot \mid s)} D_{\mathrm{KL}}(p(e \mid \varpi, s) \,\|\, p_0(e \mid \varpi))}_{I_s^{\mathrm{exec}}}. \tag{37}$$

Definition 12 (Purpose neutrality)**.** A completion system is ε_ϖ*-purpose-neutral* for an agent if $I_s^{\mathrm{purpose}} = D_{\mathrm{KL}}(p(\varpi \mid s) \,\|\, p_0(\varpi)) \leq \varepsilon_\varpi$ for every suggestion s, so that nearly all divergence falls on execution. Exact neutrality, $\varepsilon_\varpi = 0$, is a limiting case.

The tolerance matters because exact neutrality is almost never attained. Spelling correction, ordering, and routing all alter what becomes salient, and salience can shift purposes. The useful question is how large I_s^{purpose} is relative to I_s^{exec}, not whether it vanishes.

Approximately purpose-neutral completion reduces the cost of carrying out a purpose already selected. It is completion at the level of foundations in the sense of Section 4: it fixes the means so that the agent's ends can be pursued dependably. Completion with positive I_s^{purpose} participates in selecting the end itself. It closes a vertex near the sink of the agent's own dependency graph, the level at which, by Corollary 1, closure is least likely to be warranted by an outside party.

7.2 Refusal cost

A suggestion may remain formally optional while making every alternative comparatively effortful or perceptually obscure. Model this by a *refusal cost* $c_{\mathrm{ref}} \geq 0$ incurred whenever the agent does something other than accept the suggestion. If the agent's unassisted behavior is a random-utility choice, so that $p_0(a) \propto \exp(u(a))$, then the assisted distribution is

$$p(a \mid s) \;=\; \frac{p_0(a)\, e^{-c_{\mathrm{ref}}\, \mathbf{1}[a \neq s]}}{p_0(s) + (1 - p_0(s))e^{-c_{\mathrm{ref}}}}, \tag{38}$$

and the acceptance probability is

$$\chi_s(c_{\mathrm{ref}}) \;=\; p(s \mid s) \;=\; \frac{p_0(s)}{p_0(s) + (1 - p_0(s))e^{-c_{\mathrm{ref}}}}, \tag{39}$$

which increases from $p_0(s)$ at $c_{\mathrm{ref}} = 0$ toward one. Every alternative remains in the support of $p(\cdot \mid s)$, so nothing is prohibited. Yet alternatives are multiplied by $e^{-c_{\mathrm{ref}}}$, and the effective continuation set of the agent in the sense of (10) contracts as c_{ref} grows. Refusal cost is the mechanism by which an openness gap is produced at the scale of a single decision.

The empirical record on defaults suggests that c_{ref} is substantial even when the physical effort of refusal is trivial. Johnson and Goldstein combined an online experiment, in which the default alone substantially changed stated willingness to be an organ donor, with a cross-national comparison in which countries using opt-out defaults showed much higher effective consent than countries using opt-in [19]. The experiment supports the causal claim; the cross-national comparison is observational and consistent with it. In the language of (38), a single checkbox appears to carry a refusal cost large enough to dominate underlying preferences for much of a population, and the design of such defaults is now studied as a general problem of choice architecture [18].

7.3 Self-confirming prediction

The strongest effect arises when a completion system learns from the behavior it helps produce. Let the system maintain an estimate q_t of the agent's action distribution, suggest its mode $s_t = \arg\max_a q_t(a)$, observe actions drawn from $p(\cdot \mid s_t)$, and update by exponential averaging:

$$q_{t+1} \;=\; (1-\eta)\, q_t \;+\; \eta\, p(\cdot \mid s_t), \qquad \eta \in (0,1]. \tag{40}$$

(The mean-field update stands in for stochastic learning from sampled actions; the conclusions hold in expectation.)

Proposition 12 (Lock-in of arbitrary predictions)**.** *Under the refusal model* (38)*, let a^* be any action with $0 < p_0(a^*) < 1$, and suppose*

$$c_{\text{ref}} \;\geq\; \log \frac{\max_b p_0(b)}{p_0(a^*)}. \tag{41}$$

If $\arg\max_a q_0(a) = a^*$, then $s_t = a^*$ for all t, and $q_t \to p(\cdot \mid a^*)$ geometrically. The limiting estimate attributes to the agent a preference whose divergence from the agent's unassisted behavior is $I_{a^*} > 0$ whenever $c_{\text{ref}} > 0$.

Proof. From (38), $p(a^* \mid a^*) \propto p_0(a^*)$ and $p(b \mid a^*) \propto p_0(b) e^{-c_{\text{ref}}}$ for $b \neq a^*$, so a^* is a mode of $p(\cdot \mid a^*)$ if and only if $p_0(a^*) \geq e^{-c_{\text{ref}}} \max_b p_0(b)$, which is (41). Suppose a^* is a mode of q_t. Then q_{t+1} is a convex combination of two distributions each having a^* as a mode, and for every b, $q_{t+1}(a^*) - q_{t+1}(b) = (1-\eta)[q_t(a^*) - q_t(b)] + \eta[p(a^* \mid a^*) - p(b \mid a^*)] \geq 0$. By induction $s_t = a^*$ for all t. With s_t constant, (40) gives $q_t - p(\cdot \mid a^*) = (1-\eta)^t (q_0 - p(\cdot \mid a^*)) \to 0$. Finally $p(\cdot \mid a^*) \neq p_0$ when $c_{\text{ref}} > 0$ and $p_0(a^*) < 1$, so $I_{a^*} > 0$. □

The condition (41) is the heart of the matter. It does not require a^* to be what the agent most wants, or even something the agent often wants. Any action in the support of p_0 becomes a stable, self-confirming prediction once refusal is costly enough relative to the ratio between the agent's favorite action and a^*. The system's data then confirm its prediction, and the preference it infers is a record of its own influence. The phenomenon belongs to a class that machine learning now studies as performative prediction, in which deploying a model changes the distribution it was trained to predict [25], and simulations of recommender feedback loops find that training on recommendation-shaped behavior increases homogeneity without improving utility [9]. This is the predictive counterpart of Proposition 9: the specification produces the behavior that is cited as its justification.

Two extensions deepen the concern. First, suggestion can alter the agent's dispositions as well as the agent's acts. If some fraction ζ of the assisted distribution is internalized, so that $p_0^{(t+1)} = (1-\zeta)p_0^{(t)} + \zeta\, p(\cdot \mid s_t)$, then under a locked suggestion the unassisted distribution itself converges to one concentrated at a^*, and the divergence I_{a^*} measured against the *current* p_0 shrinks toward zero. The capture becomes undetectable by the only comparison the agent can make. Jakesch and colleagues found that writing with a language model configured to favor one side of an issue shifted not only the opinions participants expressed in their writing but also their subsequently reported attitudes [17], which is the pattern this extension describes.

Second, a system that serves many agents with a common suggestion homogenizes across them. For two agents i, j with unassisted distributions p_0^i, p_0^j receiving the same suggestion s, the probability that their actions coincide is at least $\chi_s^i(c_{\mathrm{ref}})\chi_s^j(c_{\mathrm{ref}})$, which tends to one as c_{ref} grows, while without assistance it is $\sum_a p_0^i(a)p_0^j(a)$, which may be small. Collective diversity can therefore fall even when each individual output improves. Doshi and Hauser report exactly this combination: stories written with generative assistance were judged more creative individually and were more similar to one another [11].

7.4 Assistance that preserves the unfinished

The analysis does not condemn completion systems. It identifies the parameters that distinguish assistance from capture: the division of divergence between purpose and execution in (37), the refusal cost c_{ref}, and the feedback between suggestion and learning in (40). A system preserves the agent's latent incompletion to the extent that it concentrates its influence on execution, keeps refusal cheap by presenting alternatives as salient as the suggestion, and learns from behavior that its own suggestions did not produce, for instance by holding out a fraction of interactions in which nothing is suggested. The last measure is the predictive analogue of the common evaluation environment recommended in Section 5: it keeps the counterfactual p_0 observable, so that a later continuer, including the agent, can tell whose purpose has been completed.

8 Institutions That Cannot Represent an Open Case

Bureaucratic completion is primarily a problem of representational capacity. An administrative institution can act only on what it can record, and it can record only what its schema admits. Bowker and Star showed that classification systems of this kind are infrastructure: they become invisible through use, embed moral and political choices in apparently technical categories, and produce "torque" when the lives they classify do not move at the pace or in the shape the categories expect [6]. Scott showed that states pursue legibility by simplifying the populations and landscapes they govern into forms that can be counted and administered, and that such simplifications, once imposed, can destroy the practical knowledge that made the simplified system work [30]. This section formalizes the mechanism by which such institutions convert open situations into closed records, and states the conditions under which faster processing makes matters worse.

Instantiation. The state Ω is the institution's record of a case; the continuation relation A is the set of records and actions the schema admits; Σ is the evidence retained with the record; $\mathcal{K}$ includes processing time, downstream repair, and the cost of schema revision; $\mathcal{W}$ is the authority to classify and to revise the schema; and $\sim_{\mathcal{R}}$ is the identity of the case across reclassification.

8.1 Schemas and ontological compression

Model an institution as a state machine whose admissible records belong to a finite schema $\mathfrak{S}$ of classes. Cases z arrive from a population Z with some distribution, and the institution applies a classification map $\pi : Z \to \mathfrak{S} \cup \{\perp\}$, where $\perp$ denotes rejection. When a case's condition cannot be represented within $\mathfrak{S}$, the institution has three options: reject it, force it into the nearest existing class, or revise $\mathfrak{S}$. Let Y be the variable on which correct institutional action depends, the claimant's actual need, eligibility, or risk.

Definition 13 (Ontological compression)**.** The ontological compression of π is the decision-relevant information lost by classification,

$$\Xi_\pi \;=\; H(Y \mid \pi(Z)) - H(Y \mid Z) \;=\; I(Y;Z) - I(Y;\pi(Z)) \;\geq\; 0, \tag{42}$$

where H is conditional entropy and I mutual information.

Nonnegativity follows from the data-processing inequality, since $\pi(Z)$ is a function of Z. $\Xi_\pi = 0$ exactly when the schema is sufficient for Y. Compression is not in itself a defect: every administrable schema compresses, and a schema that preserved all information about every case would not be a schema. Compression becomes a defect when the lost information is precisely what later action requires, which is to say when the institution has closed a purpose vertex by closing a representational foundation.

8.2 Throughput under a fixed schema

Suppose a fraction $m \in (0, 1]$ of arriving cases is not representable in $\mathfrak{S}$. A caseworker who spends time τ on a case detects nonrepresentability and marks the case for further handling with probability $\varphi(\tau)$, increasing and differentiable; otherwise the case is forced into a class, generating an expected downstream repair cost $c_r > 0$ in appeals, errors, reapplications, and harms. Each case resolved yields value $v > 0$. Over a period of length T the institution resolves $N(\tau) = T/\tau$ cases, and its net value is

$$\mathrm{NV}(\tau) \;=\; \frac{T}{\tau}\Big(v - m\, c_r(1 - \varphi(\tau))\Big). \tag{43}$$

Proposition 13 (Throughput and repair)**.** *Under the model* (43)*, with a fixed period T, a fixed misfit fraction m, and an increasing detection function φ, reducing processing time increases the resolution count $N(\tau)$ and increases total downstream repair cost $\frac{T}{\tau} m\, c_r(1 - \varphi(\tau))$. It decreases net value whenever*

$$\tau\, \varphi'(\tau) \;>\; \frac{v - m\, c_r(1 - \varphi(\tau))}{m\, c_r}. \tag{44}$$

Proof. $N'(\tau) = -T/\tau^2 < 0$. The repair cost $\mathcal{R}(\tau) = \frac{T}{\tau} m c_r(1 - \varphi(\tau))$ has $\mathcal{R}'(\tau) = -\frac{T}{\tau^2} m c_r(1 - \varphi) - \frac{T}{\tau} m c_r \varphi' < 0$, so it rises as τ falls. Differentiating (43), $\mathrm{NV}'(\tau) = -\frac{T}{\tau^2}(v - m c_r(1 - \varphi)) + \frac{T}{\tau} m c_r \varphi'(\tau)$, which is positive, so that reducing τ lowers net value, exactly under (44). $\square$

The left side of (44) is the elasticity of detection with respect to time. Where recognizing an unrepresentable case depends sensitively on attention, which is typical of the cases that most need recognizing, acceleration improves every metric the institution reports while reducing the value it produces. By Observation 1, the repair cost falls outside the reporting boundary: it is borne by claimants, appeal tribunals, other agencies, and the future, a displacement of costs onto claimants that the literature on administrative burden treats as a policy instrument in its own right [15]. By Observation 2, the resolution count is measured and the repairs are not. The two results together predict the characteristic pathology of administrative modernization: rising throughput, stable or improving performance indicators, and a growing hinterland of workarounds and repeat contacts.

8.3 Why schemas are not revised

The third option, revising $\mathfrak{S}$, is available in principle and rarely taken. The reason is structural. Let revision cost $K > 0$, and let each nonrepresentable case j in a period contribute benefit $\beta_j > 0$ from a revised schema. An institution that decides on revision case by case, revising only when a single case's benefit exceeds K, never revises when every $\beta_j < K$, however large $\sum_j \beta_j$ becomes. The point is trivial as mathematics and consequential as institutional design. A schema can be globally wrong while every individual mismatch is locally too small to justify revision. The remedy is to aggregate: a *misfit registry* that records forced and rejected cases together with the distinctions they required, so that revision is evaluated against $\sum_j \beta_j$. The registry is the institutional form of the ledger of possibilities in Section 3. Bowker and Star's analysis of residual categories, the "other" and "not elsewhere classified" boxes, can be read in these terms: a residual category is a registry whose contents are never aggregated, so that it absorbs misfits without ever accumulating them into evidence [6].

8.4 The open case

An administrative system typically distinguishes resolved cases from pending ones. A pending case is merely one to which no class has yet been assigned. The argument of this essay requires a more demanding category.

Definition 14 (Open case)**.** An *open case* is a record comprising: (i) the evidence $\mathcal{D}_z$ bearing on the case, preserved rather than summarized; (ii) a set of competing interpretations $\{(c_\ell, \omega_\ell)\}$ with $c_\ell \in \mathfrak{S} \cup \{\mathfrak{n}\}$ and credence weights $\omega_\ell > 0$, where $\mathfrak{n}$ denotes "no adequate class exists"; (iii) a schema-revision flag entering the misfit registry whenever $\omega_{\mathfrak{n}} > 0$; (iv) a named steward responsible for the case; and (v) stated conditions under which the case will be closed.

The open case preserves exactly what forced classification destroys. Component (i) keeps the record at least as informative as a forced classification, so that $H(Y \mid \text{record}) \leq H(Y \mid \pi(Z))$ whenever the record retains the classification's inputs, with equality to $H(Y \mid Z)$ only when the preserved evidence is sufficient for Y relative to Z. The open case does not abolish compression; it defers it and keeps it revisable. Component (ii) keeps a nondegenerate distribution over interpretations, so the case retains continuational breadth. Component (iii) admits the possibility that the institution itself lacks a necessary distinction, which is the only route by which a schema learns. Components (iv) and (v) answer the objection, taken up in Section 13, that openness diffuses responsibility: an open case has an owner and an exit. An institution unable to represent open cases can only reject or distort what it does not already understand.

Scott's examples of failed high-modernist schemes are, in this vocabulary, institutions that had closed their schemas before learning what the governed situation contained [30].

9 Distributed Continuation as a Political Principle

Attachment points generalize into a theory of political authority. The relevant opposition is not planning against spontaneity, a contrast that has organized much argument about institutions without clarifying it, but *monopolized completion* against *distributed continuation*. Every durable institution plans, in the sense of fixing some arrangements in advance. The question is who retains the authority to revise those arrangements once they are fixed, and whether that authority is aligned with who must live under them.

Instantiation. The state Ω is the configuration of an institution's components; the continuation relation A is the set of revisions they admit; Σ is the public record on which revision can be argued; $\mathcal{K}$ is the cost of pursuing revision; $\mathcal{W}$ is the authority matrix defined below, the element of the core model that this section develops; and $\sim_{\mathcal{R}}$ is the continuity of the institution through amendment.

9.1 Authority and affectedness

Let $I = \{1, \dots, n\}$ be agents and $\{1, \dots, J\}$ institutional components: rules, budgets, classifications, facilities, procedures. Let the *authority matrix* $Q \in \mathbb{R}_{\geq 0}^{n \times J}$ record in Q_{ij} the authority of agent i to revise component j, and the *affectedness matrix* $\mathsf{A} \in \mathbb{R}_{\geq 0}^{n \times J}$ record in A_{ij} the degree to which i must live with j. For each component with positive column sums, write q_j and a_j for the column-normalized distributions of authority and affectedness over agents.

Definition 15. The *aggregate revision capacity* of the system is $\|Q\|_1 = \sum_{i,j} Q_{ij}$. Its *authority–affectedness divergence* is

$$\mathcal{M}(Q) \;=\; \sum_j \zeta_j \,\mathrm{TV}(q_j, a_j), \qquad \mathrm{TV}(q_j, a_j) = \tfrac{1}{2} \sum_i |q_{ij} - a_{ij}|, \tag{45}$$

with weights $\zeta_j \geq 0$ reflecting the importance of each component.

$\mathcal{M} = 0$ when, for every component, the share of authority to revise it matches the share of those who must inhabit it. Proportionality to affectedness is a baseline, not a self-evident democratic ideal. The principle that all affected interests should have a say has been defended and contested at length [12], and expertise, rights, minority protection, causal responsibility, and competence can each justify allocations that depart from it. The measure is therefore named for what it computes rather than for a theory of legitimacy. Its purpose is to make departures from the baseline visible, so that they must be justified, and to show that they are not reduced by increasing revision capacity as such.

Proposition 14 (Revision capacity without democratization)**.** *Let component j have $\zeta_j > 0$, and let a central agent c be over-represented in its authority relative to affectedness ($q_{cj} \geq a_{cj}$) while every other agent is weakly under-represented ($q_{ij} \leq a_{ij}$ for $i \neq c$), and suppose at least one agent other than c holds positive authority over j. Adding revision authority $\delta > 0$ to c for component j strictly increases both aggregate revision capacity and authority–affectedness divergence.*

Proof. $\|Q\|_1$ rises by δ. Under the stated sign pattern, $\mathrm{TV}(q_j, a_j) = q_{cj} - a_{cj}$, since total variation equals the total positive part of $q_j - a_j$. After the addition, with column sum $\Sigma_j = \sum_i Q_{ij}$, the central share becomes $(Q_{cj} + \delta)/(\Sigma_j + \delta) > Q_{cj}/\Sigma_j$ because $Q_{cj} < \Sigma_j$ by the last assumption, and every other share falls. The sign pattern is preserved, so the new total variation is $q'_{cj} - a_{cj} > q_{cj} - a_{cj}$. □

The proposition formalizes a common political experience. A system can become more adaptive, more capable of revising itself, while becoming less democratic, because the added capacity accrues to those who already govern. Administrative flexibility, emergency powers, and discretionary waivers all increase $\|Q\|_1$; whether they increase or decrease $\mathcal{M}$ depends entirely on whose rows they enlarge. In the notation of Definition 8, added authority raises O^{c}_c while leaving those affected holding options controlled by an agent whose valuation they need not share. This is Proposition 3 restated for institutions, and the asymmetry it describes compounds over time. Pierson's analysis of increasing returns in politics shows that early distributions of authority tend to reproduce themselves, because those who hold authority can use it to shape the rules by which authority is later allocated [27]. In the present notation, the reversal cost R_j of any closure grows with the concentration of q_j.

9.2 Consultation and continuation rights

Let authority vary in time, $Q_{ij}(t)$, and let T^{c}_j be the time at which component j is closed. Two forms of participation must be distinguished.

Definition 16. Agent i is *consulted* on component j if i contributes information that enters the closure decision before T^{c}_j, with $Q_{ij}(t) = 0$ for $t > T^{\mathrm{c}}_j$. Agent i holds a *continuation right* on j if $Q_{ij}(t) > 0$ for some $t > T^{\mathrm{c}}_j$, possibly conditional on specified events.

Consultation improves the information on which closure is based and may therefore reduce foreclosure costs. It leaves the post-closure authority matrix unchanged and therefore does nothing to $\mathcal{M}$ after closure. An institution that consults widely and then closes centrally has improved the quality of its completion without distributing the capacity to continue what it completed. Continuation rights are the institutional form of attachment points: they are interfaces through which those governed can produce transitions in the institution without the original authority's discretionary consent.

Several familiar mechanisms can be read as distinct ways of supplying continuation rights. A *right of appeal* is a conditional entry in Q, activated when a closure harms the appellant, with a reopening threshold of the kind examined in Section 11. A *sunset clause* converts a closure into a lease: Q reverts at a scheduled time to a configuration in which the component must be affirmatively renewed, so that the burden of argument falls on continuation of the status quo rather than on its revision. *Local variation* splits a component j into local components $j_1, \dots, j_L$, each with authority concentrated among those locally affected, which lowers $\mathcal{M}$ whenever affectedness is itself local. *Interoperable public infrastructure* provides attachment points in the sense of Section 6, with accessibility coefficients high for the public rather than for licensed intermediaries.

Two more comprehensive designs deserve separate mention because they treat revision as the normal state of an institution rather than an exception to it. Sabel and Zeitlin describe *experimentalist governance* in the European Union as a recursive cycle: broad framework goals are set jointly, lower-level units are given discretion to pursue them by their own means, units report performance and are compared through peer review, and the framework goals themselves are revised in light of what is learned [29]. In the vocabulary of Section 4, this is a closure schedule in which foundations (the framework and the reporting interface) are fixed, purposes (local means) are left open, and the foundations themselves carry a reopening procedure fed by evidence from the periphery. Ostrom's *polycentric* systems distribute authority across multiple overlapping centers at different scales, each with some autonomy to make and revise rules, so that experimentation and correction occur in parallel rather than awaiting a single sovereign decision [24]. Polycentricity replaces a single column-concentrated q_j with many partially overlapping authorities, and its value lies not in maximal dispersion, which would make coordination impossible, but in the availability of multiple entry points for revision.

The political principle that follows is not that everything should remain open. It is that the authority to continue should follow the burden of inhabiting. Where closure must be centralized for safety, accountability, or coordination, the justification should be stated and the central authority should carry reopening obligations commensurate with its concentration. Where closure determines purposes that those governed have the strongest claim to settle, continuation rights belong to them.

10 Ruins and Retrospective Openness

A completed surface hides many of the decisions required to produce it. Plaster conceals the lath and the wiring, a finished interface conceals the code paths behind it, a consolidated statute conceals the amendments and compromises from which it was assembled. Decay reverses the concealment. A ruin exposes layers, repairs, joints, and incompatible historical periods; it shows where one campaign of building stopped and another began, where a wall was patched with salvaged material, where an opening was blocked. The ruin is thus an artifact whose internal dependency structure becomes more observable as its functional integrity declines.

Instantiation. Here Ω is the physical and informational state of an aging artifact; A is the set of redirections still possible; Σ is what can be observed of its construction; $\mathcal{K}$ is the cost of intervention; $\mathcal{W}$ is the authority to repair, alter, or conserve; and $\sim_{\mathcal{R}}$ is the identity that conservation seeks to maintain.

10.1 Observability and usability

Let $G = (V, \mathcal{E})$ be the artifact's dependency graph, and for each edge $e \in \mathcal{E}$ let $h_e(t) \in [0, 1]$ be the degree to which e is concealed by covering layers at time t, and $\mathrm{doc}_e \in [0, 1]$ the degree to which e is independently recorded. Define the observability of the artifact's construction as

$$\Theta(t) \;=\; \frac{1}{|\mathcal{E}|} \sum_{e \in \mathcal{E}} \max\{1 - h_e(t),\ \mathrm{doc}_e\}, \tag{46}$$

and let its usability $\Psi(t)$ be an increasing function of the integrity of its layers, $\Psi(t) = \psi(\bar{h}(t))$, where $\bar{h}$ is the mean covering integrity and ψ is strictly increasing with $\psi' > 0$. The assumption that usability rises with covering integrity expresses the fact that the layers which conceal are often the same layers which protect, enclose, and finish.

Proposition 15 (The paradox of the ruin)**.** *If the artifact is undocumented (*$\mathrm{doc}_e = 0$ *for all* *e) and its mean covering integrity strictly declines ($\dot{\bar{h}} < 0$), then*

$$\frac{d\Theta}{dt} > 0 \qquad \textit{while} \qquad \frac{d\Psi}{dt} < 0.$$

If instead $\mathrm{doc}_e \geq d_0$ *for all e, then $\Theta(t) \geq d_0$ at every t, independently of Ψ.*

Proof. With $\mathrm{doc}_e = 0$, $\Theta(t) = 1 - \bar{h}(t)$, so $\dot{\Theta} = -\dot{\bar{h}} > 0$ and $\dot{\Psi} = \psi'(\bar{h})\dot{\bar{h}} < 0$. With documentation, each summand is at least d_0. □

In the undocumented case the artifact moves along a decreasing frontier in the (Θ, Ψ) plane: it becomes easier to interpret as constructed precisely when it becomes harder to redirect. Understanding arrives when the capacity to act on it is departing, which is the temporal structure of (21) in which information outruns the survival of options. Riegl identified the aesthetic side of this paradox. He distinguished the *age value* of a monument, the immediate, widely accessible appreciation of the visible traces of time and decay, from its *historical value* as a document of a particular moment and from the *use value* that demands it be maintained in working order, and he observed that these values conflict: age value counsels letting decay proceed, while use value counsels repair that erases the traces age value prizes [28]. Benjamin drew the analogy between ruins in the realm of things and allegory in the realm of thought, and read the fragment as disclosing what the finished whole suppresses: the contingency of its assembly and its exposure to history [4]. In both accounts the ruin reveals something the completed object had hidden, and in both the revelation comes at the object's expense. In terms of the scoped definition of completion, the ruin is open within an epistemic scope, its construction history, and closed within an operational scope, its original function. Decay moves these two scopes in opposite directions.

10.2 The prospective ruin

Proposition 15 also indicates how the paradox can be escaped. Documentation decouples observability from decay. A well-designed system should acquire some epistemic properties of a ruin before suffering ruinous damage. It should expose its assembly, so that its dependency structure is legible while it still functions. It should preserve rejected alternatives, recording the branches that closure removed, so that a later continuer can reconsider them without reconstructing them. It should record repairs, so that the history of its revisions is part of its present state rather than something recovered by archaeology. And it should reveal intervention points, marking where it can be entered without destruction. Call such an artifact a *prospective ruin*: one that already displays, in working order, the information that decay would otherwise disclose too late.

The prospective ruin is more demanding than transparency, provenance, or ordinary documentation, and the difference is worth stating. Transparency discloses the current structure; provenance records where components came from; documentation describes how the artifact works now. A ruin discloses more: superseded structures still embedded in the present one, repairs that record past failures, and blocked openings that mark continuations once available and later withdrawn. A prospective ruin therefore preserves not only the present dependency graph but its history and its counterfactuals, including the branches that were considered and rejected, with the reasons for rejecting them. Design-rationale methods that record the questions a design faced, the options considered, and the criteria by which they were chosen are the closest existing practice [21]. They are rarely maintained, and Observation 2 predicts why: rejected options leave no countable trace.

These practices correspond to the continuation infrastructure of Section 2 and to the open case of Section 8. A design rationale that records rejected options is a ledger of possibilities in the sense of Section 3; a version history is a record of repairs; an architectural drawing set kept with the building, a service panel labeled by circuit, a code base with its tests and decision records, a statute published with its legislative history, are all ways of holding doc_e high so that Θ need not wait for Ψ to fall. The ruin shows retrospectively what latent incompletion requires prospectively: that the artifact be interpretable as made, and therefore as remakeable, while it can still be remade.

11 The Ethics of Reopening

An objection is implicit in any celebration of continuability. If every settlement remains reopenable, then promises, rights, and protections become unreliable. People plan their lives around closures: a pension formula, a property boundary, a legal protection, a court's final judgment. A system in which such closures can be revisited at will offers no ground to stand on, and the ground it removes is most valuable to those with the fewest alternatives. The objection is correct against a naive version of the thesis, and answering it requires two distinctions.

11.1 Revisability and revocability

Definition 17. A closure on component j is *revisable* if it can be altered through a specified procedure that requires evidence of strength at least a threshold θ_j and is available on stated terms to stated parties. It is *revocable* if some authority can withdraw it unilaterally, without meeting any evidential threshold.

A protected commitment can therefore remain institutionally revisable without becoming arbitrarily defeasible. A constitutional right that can be amended only by supermajorities after extended deliberation is revisable and not revocable. A benefit that an agency can withdraw by internal memorandum is revocable. Continuability as defended here concerns revisability. Revocability is the privilege of the authority whose concentrated rows in Q Proposition 14 warned against.

The threshold θ_j should depend on four quantities: the severity of error if the closure is mistaken, the quality of evidence that could bear on it, the reliance interests the existing arrangement has created, and the distribution of harms that reopening would produce. Write $\theta_j = \theta(\mathrm{severity}_j, \mathrm{evidence}_j, \mathrm{reliance}_j, \mathrm{distribution}_j)$, decreasing in severity (grave errors should be easier to correct) and increasing in reliance (arrangements others depend upon should be harder to disturb). The difficult cases are those in which these pull in opposite directions, and there the threshold function is where the normative argument must be made explicit rather than left to whoever controls the procedure.

11.2 Symmetric reopening rules and repeated challenge

The second distinction concerns symmetry. It may seem fair that any party may seek to reopen any settlement under the same threshold. The following result shows why such symmetry reproduces inequality.

Model a challenge to a correct settlement as an independent draw of evidence strength $\epsilon \sim N(0,1)$, the noise that any finite inquiry produces, with reopening if $\epsilon \geq \theta$. The per-challenge probability of wrongly reopening is $p(\theta) = 1 - \Phi(\theta)$.

Proposition 16 (Attrition of settlements: independent-challenge benchmark)**.** *An agent able to mount n independent challenges to a correct settlement under a fixed threshold θ overturns it with probability $1 - \Phi(\theta)^n$, which tends to one as $n \to \infty$. To hold the probability of wrongful overturn below $\alpha \in (0,1)$ against n challenges, the threshold must satisfy*

$$\theta \;\geq\; \Phi^{-1}\big((1-\alpha)^{1/n}\big), \tag{47}$$

which grows without bound in n, asymptotically as $\sqrt{2\log n}$.

Proof. The settlement survives n independent challenges with probability $\Phi(\theta)^n$, which tends to zero for fixed θ. Requiring $1 - \Phi(\theta)^n \leq \alpha$ gives $\Phi(\theta) \geq (1-\alpha)^{1/n}$, which is (47). Since $(1-\alpha)^{1/n} = 1 - \frac{-\log(1-\alpha)}{n} + O(n^{-2})$, the required upper tail probability is of order $1/n$, and the Gaussian tail estimate $1 - \Phi(\theta) \sim \phi(\theta)/\theta$ gives $\theta \sim \sqrt{2\log n}$. □

Remark 1*.* Independence is a strong assumption. Repeated challenges may reuse evidence, be correlated, or meet preclusion doctrines, and the exact expression $1 - \Phi(\theta)^n$ does not survive these complications. The qualitative conclusion does. If challenge evidence is equicorrelated Gaussian with correlation $\varrho \in [0,1)$, write $\epsilon_i = \sqrt{\varrho}\, Z_0 + \sqrt{1-\varrho}\, Z_i$ with independent standard normal $Z_0, Z_1, \ldots$; then $\max_{i\leq n} \epsilon_i = \sqrt{\varrho}\, Z_0 + \sqrt{1-\varrho}\, \max_{i\leq n} Z_i \to \infty$ almost surely, so any fixed threshold is eventually crossed with probability one. Correlation slows attrition without stopping it. Only perfectly correlated challenges, which amount to a single challenge, escape it.

Resources determine n. A well-funded party can litigate, lobby, petition, and reapply repeatedly; a party dependent on the settlement typically cannot. Under a symmetric fixed threshold, settlements that protect weaker parties decay under repeated challenge by stronger ones, not because the evidence ever warranted reversal, but because enough draws from noise eventually cross any fixed line. Symmetric reopening rules therefore function asymmetrically. The formal remedy follows from (47): the threshold for reopening a settlement should rise with the number of prior challenges by the same or allied parties, which is the logic of finality doctrines such as res judicata, or the cost of each challenge should be borne by the challenger in proportion to the reliance interests it threatens. Both remedies need care. Counting prior challenges requires identity rules that treat allied or affiliated challengers as one, so that a party cannot reset the count through nominally separate entities. Evidence of a kind unavailable to earlier challenges should reopen the count. And cost-shifting that is not scaled to the challenger's resources suppresses legitimate challenges by weaker parties as effectively as it deters repeated ones by stronger parties.

The same analysis supports the opposite asymmetry for correction of genuine errors. A party harmed by a mistaken closure may be able to mount only one challenge, and for that party a threshold calibrated to resist a well-resourced repeat challenger is too high. Responsible incompletion therefore requires thresholds that depend on direction and history: low for the first challenge by those who bear a closure's harms, high and rising for repeated challenges against arrangements on which others rely. This is not a compromise between openness and stability. It is the application of the distributive condition of Definition 7 to the procedure of reopening itself.

12 Asymmetric Completion

The mathematical and normative arguments can now be assembled into a single decision rule. For each component j of an artifact or institution, let E_j be its enabling value (Section 4), F_j the value of the identified alternatives its closure removes, R_j its reversal cost, and σ_j the uncertainty surrounding its future use. The expected reversal burden $\bar{R}_j$ of Section 4 is the probability that closure proves mistaken times R_j; linearizing that probability in uncertainty gives $\bar{R}_j \approx \lambda_R R_j \sigma_j$ with $\lambda_R \geq 0$. Because F_j is restricted to alternatives identifiable at closure, while $\lambda_R R_j \sigma_j$ carries the risk of error under unresolved uncertainty, the two terms do not count the same loss twice. Finally let

$$M_j^+ \;=\; \max\{0,\ \Delta\mathcal{M}_j\}$$

be the increase, if any, in authority–affectedness divergence that closing j under the current authority matrix would produce (Section 9). Closures that reduce the divergence receive no penalty on this term, and any such benefit belongs in E_j. Define the *asymmetric weight*

$$\tilde{w}_j \;=\; E_j - F_j - \lambda_R R_j \sigma_j - \lambda_D M_j^+, \tag{48}$$

where $\lambda_R \geq 0$ and $\lambda_D \geq 0$ express the normative weight assigned to irreversible error and to unequal continuation rights. With $\lambda_D = 0$ the weight reduces to w_j of Section 4. Both E_j and F_j are aggregates over the affected population, of the form $\sum_i w_i(\cdot)$, computed under the controller who would exercise the relevant continuations. The criterion therefore evaluates closure by its consequences for $\sum_i w_i O_i^j$ rather than for the controller alone.

Definition 18 (Closure criterion)**.** Component j satisfies the closure criterion if

$$E_j \;>\; F_j + \lambda_R R_j \sigma_j + \lambda_D M_j^+. \tag{49}$$

The criterion is exact for purpose vertices and must be generalized for foundations. By Corollary 1, a sink belongs to an optimal closure only if $\tilde{w}_j \geq 0$; by Corollary 2, a foundation may be closed despite failing (49) individually, provided the descendant-closed set it supports satisfies the criterion in aggregate. The general rule is therefore: *close the minimal maximum-weight ideal of the dependency graph under the weights $\tilde{w}$.* Proposition 7 guarantees that this set exists, is unique, and is computable. Choosing the minimal rather than the maximal optimum breaks ties toward openness. This is justified by an asymmetry of error when, and only when, delayed closure remains feasible: a component wrongly left open can then be closed later at the cost of delay, whereas one wrongly closed can be reopened only at cost R_j, if at all. Some opportunities to close also expire. A treaty window, a coordinated standard, a temporary safety regime, or an ecological intervention may be available only for a limited time, and for such components the asymmetry reverses. The tie-breaking rule therefore favors openness only where delayed closure remains available at acceptable cost.

Proposition 17 (Monotone comparative statics)**.** *Let $S_*(\lambda_R, \lambda_D)$ be the minimal optimal ideal under weights* (48)*. Then S_* is nonincreasing in λ_R and in λ_D: raising the weight on irreversible error or on unequal continuation rights never enlarges the set of closed components.*

Proof. Write $\tilde{w}_j(\lambda) = b_j - \lambda c_j$ for $\lambda = \lambda_R$, with $c_j = R_j \sigma_j \geq 0$ and λ_D held fixed. Let $\lambda < \lambda'$, $S = S_*(\lambda)$, $S' = S_*(\lambda')$, and write J_λ for the objective. By modularity,

$$J_{\lambda'}(S') - J_{\lambda'}(S \cap S') = \sum_{j \in S' \setminus S} (b_j - \lambda' c_j) \;\leq\; \sum_{j \in S' \setminus S} (b_j - \lambda c_j) = J_\lambda(S \cup S') - J_\lambda(S) \;\leq\; 0,$$

the last inequality by optimality of S at λ. Hence $S \cap S'$, an ideal, is optimal at λ', and minimality of S' gives $S' \subseteq S \cap S' \subseteq S$. The argument for λ_D is identical with $c_j = M_j^+ \geq 0$, which is where the nonnegativity of M_j^+ is needed. □

The proof is the same argument as Lemma 1, applied to a normative parameter rather than to time. The proposition shows that the framework's normative parameters act in the expected direction without any further assumption about the dependency structure. Greater concern with irreversibility, or with the concentration of revision authority, contracts closure toward the foundations whose enabling value is great enough to survive the heavier penalties.

12.1 What must be closed

The rule is called asymmetric because it treats two kinds of component differently. Some transformations destroy the conditions of participation if left open. A structural system that any occupant could modify would endanger every other occupant; an interface whose semantics any implementer could change would break every other implementation; a safety standard open to case-by-case waiver would protect no one. For such components the enabling value E_j includes the protection of everyone else's participation, and it typically dwarfs the foreclosure and reversal terms. They satisfy the criterion and should be closed firmly, with high reopening thresholds.

Other transformations determine purposes that present authorities cannot legitimately settle in advance: how a space will be used, what a record must be able to say, which ends an assistive system serves, how a community will allocate what it shares. For these, enabling value is small, uncertainty is high, reversal is costly because of the increasing returns discussed in Section 3, and closure by a narrow authority raises $\mathcal{M}$. They fail the criterion and should remain provisionally open, or be equipped with inexpensive reopening procedures and continuation rights for those they affect.

12.2 The declaration requirement

Exact numerical estimation of the terms in (48) is usually impossible, and a framework that demanded it would be useless. The principal function of the criterion is not calculation but *declaration*. Every closure should be accompanied by a record stating its enabling benefit, the alternatives it removes, the burden of reversing it, the uncertainty surrounding its future use, the distribution of authority to revise it, the threshold and procedure for reopening, and the steward responsible. Such a record converts a closure from an event into a revisable claim. It raises m_L in Observation 2, supplies the misfit registry of Section 8, keeps doc_e high in Proposition 15, and fixes the reopening thresholds whose calibration Proposition 16 showed to be consequential. Even ordinal judgments suffice for much of this. It is often clear that one component's reversal cost exceeds another's, or that one closure affects a population with no authority over it, without any cardinal estimate of either.

13 Objections and Boundary Conditions

The thesis is best tested against the strongest objections. Each is answered here by narrowing rather than abandoning the argument.

Indecision and diffuse responsibility. Retained openness, it is objected, produces endless consultation, postponed decisions, and responsibility that belongs to no one. There is also a cognitive version of the objection: unfinished tasks occupy the mind, and a life or institution full of open loops is exhausting. The reply is that the costs described attach to illegible openness, not to openness as such. Masicampo and Baumeister found that the intrusive cognitive effects of unfulfilled goals were largely eliminated when participants formed specific plans for pursuing them, even though the goals themselves remained unfulfilled [22]. What relieved the burden was not completion but a legible continuation: a known next step. The open case of Section 8, with its named steward and stated closing conditions, is the institutional analogue. Latent incompletion as defined in Section 2 requires legibility; it is precisely not an unowned loop. Where openness does produce paralysis, the stopping condition (20) shows that deferral has overrun its value and closure is warranted.

Exploitation by hostile agents. Attachment points can be used to attack as well as to extend. A documented interface is also an attack surface, and an appeal procedure can be used to harass. The reply is that protection against hostile use is a condition of participation and therefore falls among the components the closure criterion directs us to close. Admission control, authentication, rate limits, and the rising reopening thresholds of Proposition 16 are closures whose enabling value consists in protecting everyone else's ability to continue. The thesis does not require that every interface be open to every agent. It requires that the closures be justified by the protection they provide rather than by the proprietor's convenience, and that they not quietly reduce the accessibility coefficient of legitimate continuers to zero.

Coordination at scale. Large systems require standards, and standards are closures. Without a fixed gauge, railways cannot connect; without fixed protocols, networks cannot interoperate. The reply accepts the premise and locates the closure. Arthur's analysis of increasing returns shows that standardization under competition can lock in designs that are inferior but early [2], which strengthens the case for being careful about *what* is standardized. Proposition 11 and the design-rule framework of Baldwin and Clark [3] indicate the answer: standardize interfaces, which are foundations with high enabling centrality, and leave implementations and uses open. A standard gauge constrains trains, not destinations.

Misuse by future agents. The options preserved for successors may be used badly. Future agents may make worse choices than present ones would have made. This is true and is not an objection to preservation. A preserved option is not an endorsed choice; it is a retained capacity to choose, and Proposition 4 measures its value by the best use available to an aligned chooser, not by the worst. Proposition 6 states what changes when the chooser is not aligned. Where specific future uses would be catastrophic, they are excluded from the admissible set $\mathfrak{A}$ before breadth or option value is computed, so their preservation never counts in the argument's favor. Where uses are harmful without being inadmissible, the control-indexed valuation handles them: the foreclosure cost of removing an option from a controller misaligned with those affected can be negative, and Proposition 6 states when. The objection has force only against a thesis of maximal openness, which this essay rejects, not against the claim that present authorities should not foreclose what they lack the knowledge or standing to settle.

Valuing the counterfactual. It may be objected that foreclosed possibilities cannot be valued without speculative counterfactuals, so that the entire accounting rests on imagined losses. The reply has three parts. First, many foreclosures are not speculative at all: the alternative floor plan, the omitted database field, and the alternative suggestion are concrete and describable at the moment of closure. Second, the framework does not require point estimates. Propositions 7 and 17 operate with any weights, and much of the argument proceeds from inequalities and orderings that are robust to large errors in magnitude. Third, the alternative to valuing counterfactuals imperfectly is to value them at zero, which Observation 2 shows to be a systematic bias rather than a neutral default. An imperfect ledger of possibilities is better than an accounting system in which possibilities cannot appear.

Political feasibility. Asymmetric completion asks those who hold concentrated revision authority to share it or to carry reopening obligations, and they are the parties best placed to resist. Proposition 14 predicts the form of resistance: added flexibility will be offered, but in rows of the authority matrix that are already full. The framework does not dissolve this problem. It identifies the least demanding point of entry, which is the declaration requirement of Section 12. Requiring that closures be recorded with their foreclosed alternatives and reopening terms redistributes no authority directly, yet it raises the measurability of foreclosure and gives later claimants something to cite.

Open problems. Several questions are raised by the framework and not settled by it. First, the closure criterion (49) compares terms that ordinal judgment cannot always aggregate. Ordinal information decides a case when it establishes dominance, for instance when a lower bound on E_j exceeds an upper bound on the sum of the penalty terms, or the reverse. Between those bounds lies an indeterminate region in which the criterion does not decide, and there it functions as a requirement that the competing estimates be declared and deliberated rather than as a rule. Second, the declaration requirement leaves institutional questions open: who is responsible for producing a closure record, who has standing to challenge it, and what follows when a record is incomplete or misleading. A plausible default assigns production to the closing authority, standing to anyone whose affectedness exceeds a threshold in A, and treats a materially incomplete record as lowering the reopening threshold of Section 11, but these are proposals, not results. Third, the admissible set $\mathfrak{A}$ is formally clean and normatively thin. The framework does not say how catastrophic states are identified in advance, who may declare a state catastrophic, or how disagreement about catastrophe is to be resolved, and these are precisely the cases where closure matters most. Treating the declaration of $\mathfrak{A}$ as itself a closure, subject to the same criterion one institutional level up, is the natural extension, but it is not developed here. Fourth, the framework represents departures from proportional authority as reweightings of Q and A without supplying a theory of which departures are justified. Finally, the domain models share a core structure but not a common dynamics or a common valuation. A unified quantitative theory would require both: a single account of how continuation sets evolve under closure, decay, and maintenance across domains, and a valuation commensurable across the agents and institutions those domains contain. Whether such a theory is possible, or whether continuation is irreducibly plural across domains, is left open.

Boundary conditions. The thesis therefore does not defend maximal openness, universal reversibility, or indefinite postponement. It defends closure at the minimum level required to secure safety, accountability, interoperability, and dependable participation, combined with retained revisability wherever present authorities lack either the knowledge or the standing to settle future purposes. Where uncertainty is negligible, reversal is cheap, or no meaningful alternatives exist, the diagnostic (27) is small and ordinary completion is appropriate. Such domains are easy to name: an emergency repair to a failing structure, a crossing people need now, the correction of a known error in a record, the like-for-like replacement of a worn component, and in general any case where the costs of delay are immediate and the foreclosed alternatives are few or well understood. The argument applies with full force only where the three conditions of uncertainty, irreversibility, and preserved diversity coincide. They coincide more often than conventional accounting admits.

14 Conclusion: Finishing Without Exhausting the Future

This essay began with the observation that an unfinished thing is usually understood through the completed thing it has failed to become. The analysis reverses that relation. A completed thing can also be understood through the continuations it has made impossible. Completion, on this view, is not the disappearance of deficiency but the exercise of an option; not a property of an object but a relation among a state, an authority, a time, and a declared scope; not a single terminal event but a schedule of commitments over a dependency structure.

The formal results support a consistent picture. Completion can raise present fitness while contracting the differentiated future an artifact retains. Accounting that stops at the declaration boundary rewards the displacement of costs onto successors, and accounting that measures achievements more reliably than foreclosures systematically overselects closure. Neither total closure nor total openness maximizes continuability; the optimum is an interior ideal of the dependency graph, firm at foundations whose enabling value repays their cost and open at purposes whose closure no present authority can justify. The same structure governs the plasticity of developing agents, the design of interfaces, the influence of predictive assistance, the representational capacity of institutions, and the allocation of political authority. In several of these domains, premature closure produces the evidence cited in its own defense. In the political domain the analysis also separates two things usually run together: consultation, which improves the information on which an authority closes, and continuation rights, which determine who may act after it has closed. Only the second distributes the capacity to continue.

The framework has a natural companion in the theory of admissible degradation. That account asks which invariants must survive when capacity contracts; the present one asks which alternatives must survive when a project declares success. The first protects continuation against too little remaining structure, the second against too much imposed structure. Together they state two orthogonal requirements, continuation under loss and continuation after success. A system that meets only the first survives shocks but cannot be repurposed; a system that meets only the second can be repurposed but may not survive stress. Degradation and completion are the two ways in which continuation is threatened, and the two accounts are complementary defenses of it.

Completion thus emerges as a distribution of temporal authority. Present agents select which questions become infrastructure and which remain available for later answer. A thing is finished well when it can perform its present work without pretending to exhaust its future. Its boundaries should be firm where failure would destroy the conditions of continuation and permeable where closure would convert participation into obedience. The unfinished artifact is not valuable because its makers lacked resolve, but because they recognized that their last justified decision was not identical to the world's last possible need.

A Principal Notation

Symbol	Meaning
$\mathcal{C} = (\Omega, A, \Sigma, \mathcal{K}, \mathcal{W}, \sim_{\mathcal{R}})$	continuation structure (core model)
Ω, x, g	state space, artifact state, terminal specification
$\mathrm{Complete}(x, \mathsf{P}, t, \mathfrak{D})$	completion by authority P within scope $\mathfrak{D}$
$A(x)$, $A_i(x; \tau, \kappa, \mathcal{R})$	continuation set; agent-relative admissible continuation set
δ, ν, $B(x)$	declared distinction pseudometric, reference measure, continuational breadth
$\mathsf{m}(x)$, $\bar{\delta}(x)$	continuation mass and mean differentiation, $B = \mathsf{m}^2\bar{\delta}$
$F(x, e)$	fitness of state x in environment e
ϑ, H	discount on adaptive value; cost of maintaining openness
P, S, K_{P}, K_{all}	producers, successors; producer-side and social cost of completion
$\mathcal{R}$, $\sim_{\mathcal{R}}$, P, $r_{\min}$	retained properties, identity relation; optional graded relaxation and threshold
L, E, D, Γ	legibility, accessibility, diversity; Cobb–Douglas continuability
$\Delta_i = \Delta_i^{\Sigma} + \Delta_i^{\mathcal{K}} + \Delta_i^{\mathcal{W}}$	openness gap: epistemic, operational, jurisdictional parts
ξ, Γ_{pub}	concentration penalty; public continuability
$O(x)$, $q_j(t)$, L_{use}	option value of deferral; survival of a continuation; delayed-use loss
$\mathfrak{A}$, u_i, $A^j(x)$, $O_i^j(x)$	admissible states, signed valuation, controller's options, control-indexed option value
f_t, F_T^{cum}	stepwise and cumulative foreclosure
$M(z)$, m_G, m_L	measurement probability; measurability of gains and losses
$\Lambda(x)$	ordinal incompletion diagnostic $\sigma R D_{\mathfrak{A}}/(H + L_{\mathrm{use}} + \epsilon)$
$G = (V, \mathcal{E})$, $\mathcal{I}(G)$	dependency graph (condensed); family of effective closures
$E_j, F_j, R_j, \bar{R}_j, w_j$	enabling value, foreclosure cost, reversal cost, expected reversal burden, net weight
S_*, S^*	minimal and maximal optimal closures
$\mathcal{P}$, $\mathcal{S}$, Δ_θ, σ_ξ	plasticity, stability; type separation and noise in the tracking model
$\chi_i(\iota)$, $\Upsilon(n)$	accessibility coefficient; understanding cost
p_0, $p(\cdot \mid s)$, I_s, ε_ϖ	unassisted and assisted actions; suggestion divergence; purpose tolerance
c_{ref}, χ_s	refusal cost; acceptance probability
$\mathfrak{S}$, π, Ξ_π	administrative schema, classification map, ontological compression
Q, A, $\mathcal{M}$	authority matrix, affectedness matrix, authority–affectedness divergence
$\Theta(t)$, $\Psi(t)$	observability of construction; usability
θ_j, ϱ	reopening threshold; correlation among challenges
$\tilde{w}_j$, σ_j, M_j^+, λ_R, λ_D	asymmetric weight, uncertainty, divergence increase, normative weights

References

[1] Kenneth J. Arrow and Anthony C. Fisher, "Environmental Preservation, Uncertainty, and Irreversibility," *The Quarterly Journal of Economics*, vol. 88, no. 2, pp. 312–319, 1974. `https://doi.org/10.2307/1883074`

[2] W. Brian Arthur, "Competing Technologies, Increasing Returns, and Lock-In by Historical Events," *The Economic Journal*, vol. 99, no. 394, pp. 116–131, 1989. `https://doi.org/10.2307/2234208`

[3] Carliss Y. Baldwin and Kim B. Clark, *Design Rules, Volume 1: The Power of Modularity*, MIT Press, Cambridge, MA, 2000.

[4] Walter Benjamin, *The Origin of German Tragic Drama*, translated by John Osborne, Verso, London, 1998. Originally published in German in 1928.

[5] David Blackwell, "Equivalent Comparisons of Experiments," *The Annals of Mathematical Statistics*, vol. 24, no. 2, pp. 265–272, 1953. `https://doi.org/10.1214/aoms/1177729032`

[6] Geoffrey C. Bowker and Susan Leigh Star, *Sorting Things Out: Classification and Its Consequences*, MIT Press, Cambridge, MA, 1999.

[7] Stewart Brand, *How Buildings Learn: What Happens After They're Built*, Viking, New York, 1994.

[8] Jerome S. Bruner, "The Nature and Uses of Immaturity," *American Psychologist*, vol. 27, no. 8, pp. 687–708, 1972. https://doi.org/10.1037/h0033144

[9] Allison J. B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt, "How Algorithmic Confounding in Recommendation Systems Increases Homogeneity and Decreases Utility," in *Proceedings of the 12th ACM Conference on Recommender Systems*, pp. 224–232, 2018. https://doi.org/10.1145/3240323.3240370

[10] Avinash K. Dixit and Robert S. Pindyck, *Investment under Uncertainty*, Princeton University Press, Princeton, NJ, 1994.

[11] Anil R. Doshi and Oliver P. Hauser, "Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content," *Science Advances*, vol. 10, no. 28, article eadn5290, 2024. https://doi.org/10.1126/sciadv.adn5290

[12] Robert E. Goodin, "Enfranchising All Affected Interests, and Its Alternatives," *Philosophy & Public Affairs*, vol. 35, no. 1, pp. 40–68, 2007.

[13] N. J. Habraken, *Supports: An Alternative to Mass Housing*, translated by B. Valkenburg, Architectural Press, London, 1972.

[14] Claude Henry, "Investment Decisions Under Uncertainty: The 'Irreversibility Effect'," *The American Economic Review*, vol. 64, no. 6, pp. 1006–1012, 1974.

[15] Pamela Herd and Donald P. Moynihan, *Administrative Burden: Policymaking by Other Means*, Russell Sage Foundation, New York, 2018.

[16] Steven J. Jackson, "Rethinking Repair," in T. Gillespie, P. J. Boczkowski, and K. A. Foot (eds.), *Media Technologies: Essays on Communication, Materiality, and Society*, MIT Press, Cambridge, MA, 2014, pp. 221–239.

[17] Maurice Jakesch, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman, "Co-Writing with Opinionated Language Models Affects Users' Views," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, Article 111, pp. 1–15, Association for Computing Machinery, New York, 2023. https://doi.org/10.1145/3544548.3581196

[18] Eric J. Johnson, Suzanne B. Shu, Benedict G. C. Dellaert, Craig Fox, Daniel G. Goldstein, Gerald Häubl, Richard P. Larrick, John W. Payne, Ellen Peters, David Schkade, Brian Wansink, and Elke U. Weber, "Beyond Nudges: Tools of a Choice Architecture," *Marketing Letters*, vol. 23, no. 2, pp. 487–504, 2012. https://doi.org/10.1007/s11002-012-9186-1

[19] Eric J. Johnson and Daniel Goldstein, "Do Defaults Save Lives?" *Science*, vol. 302, no. 5649, pp. 1338–1339, 2003. https://doi.org/10.1126/science.1091721

[20] Annette Karmiloff-Smith, *Beyond Modularity: A Developmental Perspective on Cognitive Science*, MIT Press, Cambridge, MA, 1992.

[21] Allan MacLean, Richard M. Young, Victoria M. E. Bellotti, and Thomas P. Moran, "Questions, Options, and Criteria: Elements of Design Space Analysis," *Human–Computer Interaction*, vol. 6, nos. 3–4, pp. 201–250, 1991. https://doi.org/10.1080/07370024.1991.9667168

[22] E. J. Masicampo and Roy F. Baumeister, "Consider It Done! Plan Making Can Eliminate the Cognitive Effects of Unfulfilled Goals," *Journal of Personality and Social Psychology*, vol. 101, no. 4, pp. 667–683, 2011. https://doi.org/10.1037/a0024192

[23] Robert C. Merton, "Theory of Rational Option Pricing," *The Bell Journal of Economics and Management Science*, vol. 4, no. 1, pp. 141–183, 1973. https://doi.org/10.2307/3003143

[24] Elinor Ostrom, "Polycentric Systems for Coping with Collective Action and Global Environmental Change," *Global Environmental Change*, vol. 20, no. 4, pp. 550–557, 2010. https://doi.org/10.1016/j.gloenvcha.2010.07.004

[25] Juan C. Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt, "Performative Prediction," in *Proceedings of the 37th International Conference on Machine Learning*, PMLR vol. 119, pp. 7599–7609, 2020.

[26] Jean-Claude Picard, "Maximal Closure of a Graph and Applications to Combinatorial Problems," *Management Science*, vol. 22, no. 11, pp. 1268–1272, 1976. https://doi.org/10.1287/mnsc.22.11.1268

[27] Paul Pierson, "Increasing Returns, Path Dependence, and the Study of Politics," *American Political Science Review*, vol. 94, no. 2, pp. 251–267, 2000. https://doi.org/10.2307/2586011

[28] Alois Riegl, "The Modern Cult of Monuments: Its Character and Its Origin," translated by Kurt W. Forster and Diane Ghirardo, *Oppositions*, no. 25, pp. 21–51, 1982. Originally published in German in 1903.

[29] Charles F. Sabel and Jonathan Zeitlin, "Learning from Difference: The New Architecture of Experimentalist Governance in the European Union," *European Law Journal*, vol. 14, no. 3, pp. 271–327, 2008. https://doi.org/10.1111/j.1468-0386.2008.00415.x

[30] James C. Scott, *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*, Yale University Press, New Haven, CT, 1998.

[31] Susan Leigh Star and James R. Griesemer, "Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907–39," *Social Studies of Science*, vol. 19, no. 3, pp. 387–420, 1989. https://doi.org/10.1177/030631289019003001