

Exaptation Under Delayed Evaluation:

Selection-Pressure Management as a Mechanism of Creative Intelligence

Flyxion

January 27, 2026

Abstract

This essay develops a mechanism-level account of creative intelligence grounded in two coupled ideas: exaptation as the generative operator that repurposes existing elements for novel functions, and delayed evaluation as the enabling buffer that preserves function-neutrality long enough for exaptation to occur. The central claim is that exaptation is not primarily limited by the availability of parts, ideas, or skills, but by selection pressure that prematurely assigns function and imposes evaluative collapse. We formalize creative intelligence as a process that maintains reservoirs of function-neutral elements, actively protects those reservoirs from premature assessment, continuously scans for repurposing opportunities, and tests candidate repurposings against real constraints. The essay provides a mathematical framework in which evaluation timing appears as a control parameter that governs the measure of reachable functions and the expected yield of viable repurposings. We prove that, under broad conditions, earlier and harsher evaluation reduces expected creative yield by shrinking the set of admissible repurposing trajectories and by inducing path-dependent lock-in. We then interpret educational assessment regimes, platform engagement metrics, and institutional optimization mandates as selection mechanisms that systematically destroy the exaptation window, producing environments rich in materials yet poor in generativity. The result is a unified explanation of why repair, craft, and deep learning frequently require slack, dormancy, and apparent inefficiency, and why systems optimized for throughput and immediate legibility tend to suppress innovation.

1 Introduction

Creative intelligence is often described either as the production of novelty or as the capacity to explore large spaces of possibilities. Such descriptions capture important surface features but frequently mislocate the core mechanism. Novelty is not a primitive operation; it is an outcome of transformations performed on existing structures. Likewise, abundant possibility is not sufficient for creativity; possibility must remain available long enough for repurposing trajectories to form. The present essay advances a more specific claim: creative intelligence is best understood as exaptation under actively protected delayed evaluation.

Exaptation, in evolutionary theory, names the reassignment of function: a trait that emerged or persisted under one set of selective pressures is later repurposed for a different role. Feathers, once associated with thermoregulation, can later participate in flight; anatomical structures can migrate across functional regimes; biological “waste” can become regulatory substrate. The key point is not the historical examples but the operator itself: function is not fixed by design intent, and innovation proceeds by reinterpreting existing elements as candidates for new tasks.

Delayed evaluation names the enabling condition that makes this operator effective. Exaptation requires an interval in which elements can remain function-neutral with respect to the new target, so that repurposing hypotheses can be generated, explored, and tested. If evaluation is immediate, exaptation is suppressed before it begins: elements are forced into narrow roles, and trajectories that would have become viable under later constraints are eliminated early. In practice, this occurs whenever systems impose rapid, externally defined success criteria, demand immediate demonstrable performance, and treat slack as waste. The consequence is not merely that creativity is discouraged, but that the mechanism of creativity is structurally prevented.

The essay is written in four movements. First, we provide background and terminology that separates exaptation from related notions such as “loose parts,” bricolage, and transfer. Second, we formalize exaptation in a mathematical lexicon that makes evaluation timing explicit. Third, we prove several results that make precise the intuition that premature evaluation shrinks the reachable function space and reduces expected creative yield. Fourth, we interpret institutions and platforms as selection systems, showing how common evaluation regimes predictably suppress exaptation windows and thereby degrade creative intelligence.

Throughout, we avoid appeals to personal temperament, genius, or narrative charisma. The mechanism under study is not a personality trait; it is a structural relationship between repurposing operators and selection dynamics. The relevant question is not whether agents “have imagination” but whether their environments preserve the degrees of freedom required for repurposing to complete.

2 Background: Exaptation, Function-Neutrality, and the Timing of Selection

Exaptation is often introduced as a historical label for evolutionary phenomena, but the present argument treats it as a general operator on structured systems. The operator can be described

informally as follows. An element that is embedded in a system with some role, affordance profile, or causal participation is re-specified as serving a different role in a different context, without requiring the element to have been designed for the new role. The operator is therefore a mapping from an element-and-context pair to a new element-and-context pair in which the element is conserved but the function assignment changes.

This differs from mere reuse. Reuse can be trivial; exaptation is nontrivial because it changes the functional interpretation. It also differs from “loose parts” accounts that focus on variable density. Loose parts characterize the availability of manipulable elements, but they do not specify the process by which those elements become newly functional. Exaptation supplies that process: it is the rule by which function is reassigned.

The central enabling concept is function-neutrality. An element is function-neutral relative to a target if, prior to evaluation, it is not pre-committed to a single designed purpose under a regime that punishes deviation. Function-neutrality is not the same as being valueless. It means that the element is not yet collapsed into a narrow function class by external assessment. In human contexts, scrap materials, drafts, notebooks, dormant repositories, and informal prototypes are often function-neutral because they escape immediate evaluation. In institutional contexts, standardized curricula, optimized tooling, and performance-scored outputs tend to be function-committed because they encode expectations about what the element is for.

Delayed evaluation is the mechanism that maintains function-neutrality long enough for exaptation to act. Evaluation is any mapping from trajectories or states to acceptability, reward, or elimination. When evaluation occurs early, and especially when it is sharp, it prunes trajectories before they can traverse the intermediate states needed for repurposing. When evaluation is delayed, trajectories can explore intermediate states that are not immediately legible as successful but later become decisive stepping stones.

These ideas can be framed as a general hypothesis.

Definition 1 (Exaptation Operator). Let x be an element in a system and let \mathcal{C} denote a context space. A function assignment is a map $f : \mathcal{C} \times \mathcal{X} \rightarrow \mathcal{F}$ that assigns to an element $x \in \mathcal{X}$ in context $c \in \mathcal{C}$ a function label $f(c, x) \in \mathcal{F}$. An exaptation operator is any transformation E that produces a new context and function assignment for the *same* element:

$$\mathsf{E} : (c, x, f) \mapsto (c', x, f') \quad \text{with} \quad f'(c', x) \neq f(c, x)$$

such that the new function assignment is realized by feasible dynamics of the system rather than by redesign of the element.

The phrase “realized by feasible dynamics” is the crucial constraint. It distinguishes exaptation from arbitrary relabeling. A repurposing is only meaningful if the element can, under some reachable transformation of context and configuration, actually contribute to the new function.

Definition 2 (Evaluation and Evaluation Timing). Let Γ be a space of trajectories $\gamma : [0, T] \rightarrow \mathcal{S}$ through a state space \mathcal{S} . An evaluation functional is a map $\mathcal{V} : \Gamma \rightarrow \mathbb{R} \cup \{\infty\}$ that assigns cost (or

negative utility) to a trajectory. An evaluation regime with timing parameter $\tau \in [0, T]$ induces a constrained trajectory set

$$\Gamma_\tau := \{\gamma \in \Gamma : \mathcal{V}_{\leq \tau}(\gamma) < \infty\},$$

where $\mathcal{V}_{\leq \tau}$ denotes the restriction of evaluation to the prefix $[0, \tau]$. Smaller τ corresponds to earlier evaluation. Sharper regimes correspond to larger regions of Γ mapped to ∞ (elimination).

This definition treats evaluation as trajectory pruning. In many contexts, evaluation is not literally infinite cost, but it functions as elimination: students abandon approaches that fail early; institutions defund projects without quick outcomes; platforms suppress content that does not immediately engage. The mathematical abstraction of elimination captures the structural role of evaluation even when it is implemented as reduced attention, funding, or opportunity.

The main thesis can now be stated without metaphor.

Proposition 1 (Mechanism Thesis). *Creative intelligence, understood as the production of viable novel functions from existing elements, is governed by two coupled determinants: the availability of function-neutral reservoirs (inputs to exaptation), and the evaluation timing that preserves the reachability of repurposing trajectories. Exaptation is the generative operator; delayed evaluation is the enabling buffer.*

The remainder of the essay formalizes and defends this thesis.

3 A Problem-Space Formalism for Repurposing

We now define a problem-space lexicon that makes repurposing explicit. The aim is not to reduce creativity to a single scalar, but to isolate the structural dependencies that govern whether repurposing can occur at all.

Definition 3 (Repurposing Problem Space). A repurposing problem space is a sextuple

$$\mathcal{P} = \langle \mathcal{S}, \mathcal{O}, \mathcal{C}, \mathcal{E}, \mathcal{H}, \mathcal{A} \rangle,$$

where \mathcal{S} is a state space, \mathcal{O} is a set of operators (actions or transformations), $\mathcal{C} \subseteq \mathcal{S} \times \mathcal{O}$ is a constraint relation, $\mathcal{E} : \mathcal{S} \rightarrow \mathbb{R}$ is an evaluation of states, $\mathcal{H} \in \mathbb{N}$ is a horizon parameter, and \mathcal{A} is a function-assignment space that maps elements to roles under contexts. A policy is a map π selecting operators as a function of state and assignment.

The inclusion of \mathcal{A} is what distinguishes repurposing from ordinary planning. In a standard planning problem, the set of available operators is fixed and evaluation is over states. In repurposing problems, the functional interpretation of elements is itself part of the search, and the operator set may depend on which interpretation is currently active. Informally, once an element is reinterpreted as having a new role, new actions become meaningful, while others become irrelevant.

To model this dependence, let $\alpha \in \mathcal{A}$ denote a current function assignment, and define an operator availability map

$$\text{Avail} : \mathcal{S} \times \mathcal{A} \rightarrow 2^{\mathcal{O}},$$

so that the agent at state s under assignment α may only choose $o \in \text{Avail}(s, \alpha)$ subject to constraints.

A repurposing event is then a transition in assignment space.

Definition 4 (Exaptation Transition). An exaptation transition is an operator $o_{\text{ex}} \in \mathcal{O}$ such that for some states $s, s' \in \mathcal{S}$ and assignments $\alpha, \alpha' \in \mathcal{A}$,

$$(s, \alpha) \xrightarrow{o_{\text{ex}}} (s', \alpha') \quad \text{with} \quad \alpha' \neq \alpha,$$

and the transition respects constraints: $(s, o_{\text{ex}}) \notin \mathcal{C}$.

Evaluation timing enters by specifying when assignments are penalized. To capture the idea of premature function assignment, we define a selection functional that penalizes deviations from a preferred assignment class early in time.

Definition 5 (Selection Pressure Functional). Let time be discrete $t = 0, 1, \dots, H$. A selection pressure functional is a sequence of penalties $\lambda_t \geq 0$ and a preferred assignment set $\mathcal{A}^* \subseteq \mathcal{A}$ such that trajectories incurring assignment $\alpha_t \notin \mathcal{A}^*$ are penalized by λ_t . The cumulative selection cost of a trajectory $(s_t, \alpha_t)_{t=0}^H$ is

$$J_{\text{sel}} = \sum_{t=0}^H \lambda_t \mathbf{1}\{\alpha_t \notin \mathcal{A}^*\}.$$

Early evaluation corresponds to large λ_t for small t .

This formalism captures a broad family of real systems. In an educational setting, \mathcal{A}^* may be “acceptable methods” and large early λ_t corresponds to grading that punishes exploration. In a platform setting, \mathcal{A}^* may be “content that performs” and large early λ_t corresponds to engagement-based suppression. In a funding setting, \mathcal{A}^* may be “immediately legible deliverables” and large early λ_t corresponds to milestone gating.

The next section proves that increasing early selection pressure decreases the measure of reachable repurposing trajectories under mild assumptions.

4 Core Results: Premature Evaluation Shrinks the Exaptation Window

We now prove results that formalize the intuition that delayed evaluation is not merely motivational but mechanistic. The key phenomenon is reachability. Exaptation often requires traversing intermediate assignments that do not belong to the preferred set \mathcal{A}^* . If those assignments are eliminated early, the repurposing event becomes unreachable even if it would have produced superior outcomes later.

We begin with a reachability notion.

Definition 6 (Exaptation Reachability). Fix an initial pair (s_0, α_0) . An assignment α' is exaptation-reachable within horizon H if there exists a trajectory $(s_t, \alpha_t)_{t=0}^H$ respecting constraints and operator availability such that $\alpha_H = \alpha'$. Let $\mathcal{R}_H(\lambda)$ denote the set of assignments reachable under selection penalties $\{\lambda_t\}$ when the policy is constrained to keep expected selection cost below a budget B .

We adopt a mild abstraction: policies that incur too much selection cost are effectively suppressed. This models elimination, discouragement, defunding, or invisibility.

Theorem 1 (Early Selection Shrinks Reachable Assignment Sets). *Assume the agent's feasible trajectories under constraints are nonempty and that exaptation transitions require passing through at least one intermediate assignment outside \mathcal{A}^* . Consider two selection schedules λ and $\tilde{\lambda}$ such that $\tilde{\lambda}_t \geq \lambda_t$ for all $t \leq \tau$ and $\tilde{\lambda}_t = \lambda_t$ for all $t > \tau$. Then, for any fixed budget B , the reachable assignment set satisfies*

$$\mathcal{R}_H(\tilde{\lambda}) \subseteq \mathcal{R}_H(\lambda).$$

Moreover, if there exists at least one exaptation-reachable α' whose every realizing trajectory passes through an intermediate $\alpha_t \notin \mathcal{A}^*$ with $t \leq \tau$, then the inclusion is strict.

Proof. Let $\Pi(\lambda)$ denote the set of policies whose induced trajectory distribution satisfies $\mathbb{E}[J_{\text{sel}}] \leq B$ under penalties λ . If $\tilde{\lambda}_t \geq \lambda_t$ for $t \leq \tau$, then for any fixed policy π and any trajectory, the selection cost under $\tilde{\lambda}$ is at least that under λ . Hence $\mathbb{E}_\pi[J_{\text{sel}}(\tilde{\lambda})] \geq \mathbb{E}_\pi[J_{\text{sel}}(\lambda)]$. Therefore $\Pi(\tilde{\lambda}) \subseteq \Pi(\lambda)$, since any policy feasible under stricter early penalties is feasible under weaker ones, but not conversely.

Reachable assignment sets are unions over feasible policies of the assignments they can reach within horizon H . Since the feasible policy set shrinks, the reachable assignment set cannot expand, giving $\mathcal{R}_H(\tilde{\lambda}) \subseteq \mathcal{R}_H(\lambda)$.

For strictness, suppose there exists α' such that every trajectory realizing α' passes through some $t \leq \tau$ with $\alpha_t \notin \mathcal{A}^*$. Then raising penalties on those early deviations increases J_{sel} for every realizing trajectory by at least a positive amount. For sufficiently tight budget B , there exist schedules where policies that could realize α' under λ become infeasible under $\tilde{\lambda}$, removing α' from the reachable set. Hence inclusion is strict. \square

This theorem states the mechanism in mathematical form: early selection pressure reduces the reachability of exaptation states. The result is independent of the content of \mathcal{A}^* and depends only on the existence of intermediate deviations required for repurposing.

We next connect reachability to expected creative yield by introducing a notion of viable repurposings.

Definition 7 (Viable Repurposing Yield). Let $\mathcal{G} \subseteq \mathcal{A}$ be a set of “good” assignments corresponding to viable new functions under real constraints. For a selection schedule λ , define the maximal attainable yield as

$$Y_H(\lambda) := \sup_{\pi \in \Pi(\lambda)} \mathbb{P}_\pi(\alpha_H \in \mathcal{G}).$$

Corollary 1 (Early Selection Weakly Decreases Maximal Yield). *Under the hypotheses of the theorem, if $\mathcal{G} \subseteq \mathcal{R}_H(\lambda)$, then*

$$Y_H(\tilde{\lambda}) \leq Y_H(\lambda).$$

If, moreover, \mathcal{G} contains an assignment excluded by strict shrinkage, then the inequality is strict.

Proof. Since $\Pi(\tilde{\lambda}) \subseteq \Pi(\lambda)$, the supremum over a smaller set cannot exceed the supremum over a larger set. Strictness follows if the removed reachable assignment lies in \mathcal{G} and was attainable with positive probability under some feasible policy in $\Pi(\lambda)$ but not under any in $\Pi(\tilde{\lambda})$. \square

These results formalize the claim that delayed evaluation is enabling: it preserves reachability and thereby preserves attainable yield.

5 Scope and Next Steps

The preceding sections have established the formal vocabulary and proved the first core result: early selection pressure shrinks the reachable assignment set and therefore weakly decreases maximal viable repurposing yield. The next sections will deepen the framework by modeling reservoirs of function-neutral elements as distributions over assignments, relating slack to entropy and option value, and proving additional results about lock-in, path dependence, and the trade-off between short-horizon efficiency and long-horizon innovation. We will then apply the framework to educational assessment, institutional optimization mandates, and platform engagement metrics as concrete selection regimes that predictably destroy exaptation windows.

6 Reservoirs, Slack, and Option Value

The preceding analysis treated exaptation reachability primarily as a function of evaluation timing. We now extend the framework to account explicitly for reservoirs of function-neutral elements and the role of slack. Intuitively, slack is often dismissed as inefficiency. In the present framework, slack acquires a precise functional interpretation: it preserves option value in assignment space.

Let a reservoir be modeled as a probability distribution over assignments rather than a single committed function.

Definition 8 (Function-Neutral Reservoir). A function-neutral reservoir is a probability measure μ on the assignment space \mathcal{A} such that μ has support on a nontrivial subset of assignments not contained in \mathcal{A}^* . The entropy

$$H(\mu) := - \sum_{\alpha \in \mathcal{A}} \mu(\alpha) \log \mu(\alpha)$$

measures assignment diversity within the reservoir.

A highly optimized system corresponds to a degenerate reservoir concentrated on a narrow assignment set, often a singleton. A slack-rich system corresponds to a broad distribution with

high entropy. The importance of entropy here is not informational in the Shannon sense alone, but combinatorial: higher entropy implies more distinct repurposing hypotheses can be generated without acquiring new elements.

We now formalize the notion of option value.

Definition 9 (Exaptation Option Value). Given a reservoir μ and a horizon H , define the option value of the reservoir under selection schedule λ as

$$\Omega(\mu, \lambda, H) := \mathbb{E}_{\alpha_0 \sim \mu} [Y_H(\lambda | \alpha_0)],$$

where $Y_H(\lambda | \alpha_0)$ is the maximal viable repurposing yield starting from initial assignment α_0 .

Option value captures the expected creative yield of a reservoir before committing to any particular function. Slack is valuable precisely insofar as it increases Ω .

Theorem 2 (Slack Preserves Option Value Under Delayed Evaluation). *Fix horizon H and selection schedule λ . If μ_1 and μ_2 are reservoirs with $\text{supp}(\mu_1) \subseteq \text{supp}(\mu_2)$ and $H(\mu_1) < H(\mu_2)$, then for any selection schedule λ with sufficiently delayed evaluation,*

$$\Omega(\mu_1, \lambda, H) \leq \Omega(\mu_2, \lambda, H).$$

Proof. Under delayed evaluation, reachable assignment sets are determined primarily by structural constraints rather than early penalties. Since $\text{supp}(\mu_1) \subseteq \text{supp}(\mu_2)$, any initial assignment reachable from μ_1 is also reachable from μ_2 . The larger support of μ_2 admits additional initial assignments that may lie on repurposing trajectories leading to \mathcal{G} . Since $Y_H(\lambda | \alpha_0) \geq 0$ for all α_0 , linearity of expectation yields the inequality. Higher entropy increases the measure of starting points whose reachable sets intersect \mathcal{G} . \square

This theorem formalizes why slack appears wasteful under short horizons yet indispensable under long horizons. Under early evaluation, reservoirs collapse rapidly to low-entropy distributions, and option value is destroyed. Under delayed evaluation, entropy is preserved long enough for exaptation scanning to exploit it.

7 Path Dependence and Lock-In

We now address a second failure mode induced by premature selection: path dependence. Even when evaluation does not immediately eliminate trajectories, it can bias exploration toward locally legible assignments that foreclose later repurposings.

Definition 10 (Path Dependence). A repurposing process exhibits path dependence if there exist assignments α_A, α_B and a target $\alpha' \in \mathcal{G}$ such that α' is reachable from α_A but not from α_B , and early selection pressure increases the probability of committing to α_B over α_A .

Path dependence is particularly damaging when α_B appears superior under short-horizon evaluation but is inferior under long-horizon constraints.

Proposition 2 (Early Evaluation Induces Lock-In). *Assume there exist assignments α_A, α_B with $\mathcal{E}(\alpha_B) < \mathcal{E}(\alpha_A)$ at early times but such that α_A lies on a trajectory to some $\alpha' \in \mathcal{G}$ while α_B does not. Then sufficiently strong early selection pressure produces lock-in to α_B and eliminates access to α' .*

Proof. Strong early selection penalizes deviations from low \mathcal{E} assignments. Since $\mathcal{E}(\alpha_B) < \mathcal{E}(\alpha_A)$ initially, policies that explore α_A incur higher early cost and are suppressed. Once the system commits to α_B , reachability of α' is lost by assumption. Thus early evaluation induces irreversible lock-in. \square

This proposition captures a ubiquitous phenomenon: systems that optimize too early become brittle. The brittleness is not accidental; it is a direct consequence of collapsing exploration before repurposing paths can be traversed.

8 Institutional Selection Regimes

The formal results above can now be interpreted at the institutional level. Educational systems, funding agencies, and platforms implement selection schedules that correspond to particular λ_t profiles. Early grading, milestone-based funding, and engagement-driven ranking all correspond to large early penalties. In contrast, notebooks, ungraded exploration, archival repositories, and dormant projects correspond to delayed evaluation regimes.

The analysis predicts a characteristic pathology: institutions may supply abundant materials, courses, tools, or content, yet still suppress creativity because they collapse function-neutrality through early assessment. From the perspective of the model, this is not a paradox. Material abundance without delayed evaluation produces low option value.

Moreover, the model clarifies why reforms that merely add “creative activities” often fail. If evaluation timing is unchanged, adding loose parts increases \mathcal{S} but does not increase $\mathcal{R}_H(\lambda)$ or $\Omega(\mu, \lambda, H)$. The exaptation window remains closed.

9 Creativity as Selection-Pressure Management

The formalism developed here reframes creativity as a control problem. The central variable is not novelty generation but the timing and intensity of selection. Agents capable of creative intelligence are those embedded in environments that actively manage selection pressure, preserving slack and function-neutrality until real constraints can meaningfully discriminate among repurposings.

This reframing has several implications. First, it explains why repair, maintenance, and craft are reliable sources of innovation: they operate under delayed evaluation imposed by physical constraints rather than external metrics. Second, it explains why high-throughput platforms and

hyper-optimized institutions systematically underproduce genuine novelty despite enormous activity. Third, it suggests that ethical and governance questions about who controls evaluation timing are inseparable from questions about creative capacity.

10 Mechanism Overview: Exaptation and Delayed Evaluation

This paper advances a mechanism-level account of creative intelligence grounded in two coupled components: exaptation as the generative operator and delayed evaluation as the enabling buffer. Rather than treating creativity as an emergent or irreducible faculty, the account formalizes it as a predictable outcome of how selection pressure is applied over trajectories in assignment space.

By modeling evaluation timing as a form of selection pressure, we show that premature evaluation contracts reachable assignment sets, reduces option value, and induces early lock-in. Conversely, environments that preserve slack, redundancy, and apparent inefficiency maintain larger regions of possibility space in which repurposing can occur. From this perspective, what are often dismissed as wasteful or undisciplined practices function as structural supports for adaptation under uncertainty.

The framework unifies phenomena that are typically analyzed in isolation, including biological adaptation, learning and transfer, repair and maintenance, institutional design, and cultural persistence. In each domain, systems optimized for immediate legibility and performance exhibit reduced long-horizon generativity, while systems that delay evaluation preserve the conditions under which new functions can emerge.

The sections that follow elaborate this mechanism across scales and contexts. We develop formal results characterizing entropy loss, irreversibility, and metric saturation; examine biological and technological case studies where selection pressure is mismanaged; and analyze the governance implications of evaluation timing. Together, these analyses aim to show that the central challenge facing creative systems is not the absence of novelty, but the premature foreclosure of possibility.

11 Exaptation and Learning Theory: Beyond Transfer and Generalization

Standard theories of learning in cognitive science and machine learning are typically organized around the notions of generalization and transfer. In these frameworks, an agent is trained on a task or family of tasks drawn from a known or assumed distribution, and learning is successful insofar as performance improves on unseen instances drawn from the same or a closely related distribution. Transfer learning extends this idea by allowing representations learned for task *A* to be reused for task *B*, provided that the tasks share sufficient structural similarity. While these models capture important phenomena, they implicitly assume that the space of future tasks is at least partially known in advance.

Exaptation-based learning operates under a fundamentally different assumption. Rather than optimizing representations for performance on a known task family, exaptation presupposes ignorance

of future tasks and treats this ignorance as irreducible. The objective is not to minimize expected loss over a predefined distribution, but to preserve representational and functional degrees of freedom so that elements can later be reassigned to unforeseen purposes. From this perspective, transfer learning and exaptation are not points on a continuum but distinct optimization problems with incompatible objectives.

This distinction can be formalized by contrasting two learning criteria. Let \mathcal{T} denote a space of tasks and let \mathcal{D} be a probability distribution over \mathcal{T} . In transfer learning, one seeks a representation r that minimizes expected loss

$$\mathbb{E}_{T \sim \mathcal{D}}[L(T, r)],$$

possibly subject to regularization. The optimal representation under this criterion is one that compresses information in a way that is maximally useful for the task distribution \mathcal{D} . As \mathcal{D} becomes more concentrated, the representation becomes increasingly specialized.

By contrast, exaptation-based learning assumes that future tasks are not drawn from a known distribution and may lie arbitrarily far from past tasks. A natural objective in this setting is to maximize option value over tasks, defined as the expected maximum attainable performance after task revelation:

$$\mathbb{E}_T \left[\sup_{\pi \in \Pi(r)} \text{Perf}(T, \pi) \right],$$

where $\Pi(r)$ denotes the set of policies enabled by representation r . This objective rewards representations that preserve multiple latent affordances rather than those that perform optimally on any single anticipated task.

The tension between these objectives can be made precise using entropy. Let $Z(r)$ denote the set of task-relevant features preserved by representation r , and let $H(Z(r))$ denote their entropy. Under mild assumptions, minimizing expected loss over a fixed task distribution \mathcal{D} induces a contraction in $H(Z(r))$, since features irrelevant to \mathcal{D} are suppressed. This contraction improves short-horizon performance but reduces the dimensionality of the space in which future exaptations can occur.

Proposition 3. *Let r^* be an optimal representation for minimizing expected loss over a task distribution \mathcal{D} . If \mathcal{D} has finite support and excludes tasks requiring feature z , then z is eliminated from $Z(r^*)$ under sufficient optimization pressure. Consequently, the option value of r^* with respect to tasks requiring z is zero.*

Proof. Under standard assumptions in representation learning, features that do not contribute to reducing expected loss are penalized either explicitly through regularization or implicitly through gradient descent dynamics. If z is irrelevant for all tasks in the support of \mathcal{D} , its contribution to expected loss is null. Hence any representation retaining z is strictly dominated by one that discards it, yielding a lower-complexity solution. Once discarded, z cannot be recovered without retraining or external intervention, and thus cannot support exaptation to tasks requiring z . \square

This result illustrates a core asymmetry. Optimization for transfer learning is conservative with respect to known tasks but destructive with respect to unknown ones. Exaptation-based learning,

by contrast, treats unknown tasks as first-class constraints and therefore resists early compression of representational space.

Educational systems often conflate these objectives by assuming that mastery of a curriculum implies readiness for novelty. The present analysis shows that this inference is unwarranted. Curricula optimized for assessment performance induce representational collapse analogous to overfitting in machine learning. They may produce excellent transfer within narrow domains while simultaneously eliminating the very slack required for exaptation across domains.

The implication is not that transfer learning is misguided, but that it is insufficient as a general account of creative intelligence. Where future task structure is uncertain or adversarial, preserving function-neutral representations under delayed evaluation is not a luxury but a necessity. Exaptation-based learning therefore requires institutional arrangements that explicitly protect representational entropy from premature optimization, even at the cost of short-term inefficiency.

12 Repair, Maintenance, and Degradation as Sites of Exaptation

The formal framework developed above clarifies why repair and maintenance occupy a privileged position in the ecology of creative intelligence. Unlike design-from-scratch problems, repair problems arise in contexts where failure has already occurred, constraints are partially revealed, and objectives are often underspecified or evolving. These features force a natural delay in evaluation and thereby create conditions under which exaptation can operate effectively.

To formalize this intuition, consider repair as a problem defined not by an explicit target state but by the restoration of viability under uncertain failure modes. Let \mathcal{S} denote the state space of a system, and let $\mathcal{F} \subset \mathcal{S}$ denote a failure set. A repair task is initiated when the system enters \mathcal{F} , but the precise location and structure of the failure may be unknown. Evaluation is therefore indirect: success is measured by the reestablishment of acceptable behavior, not by adherence to a predefined blueprint.

This uncertainty induces delayed evaluation by necessity. In contrast to design problems, where evaluation can be specified in advance, repair problems require exploratory interaction with the system to discover which constraints are binding. During this exploratory phase, function-neutral elements retain their status precisely because no authoritative evaluation can yet be applied. Premature commitment to a particular diagnosis or solution often worsens outcomes, a phenomenon well known in engineering practice.

We can model repair as search on a dynamically revealed constraint graph. Let \mathcal{C}_t denote the set of constraints known at time t , with \mathcal{C}_0 minimal and \mathcal{C}_t expanding as exploration proceeds. Operators that appear viable under \mathcal{C}_0 may later be ruled out, while others become feasible only after intermediate modifications. Crucially, exaptation transitions often occur during this exploratory phase, when elements originally intended for unrelated purposes are discovered to satisfy emergent constraints.

Proposition 4. *In repair problems with dynamically revealed constraints, any policy that commits*

to a fixed function assignment prior to full constraint revelation is strictly dominated by a policy that preserves assignment flexibility until constraint revelation stabilizes.

Proof. Let α denote an early committed assignment and let α' denote an alternative assignment that becomes viable only after constraint revelation at time $t > 0$. A policy that commits to α cannot access α' once commitment is enforced, whereas a flexible policy can select α' after constraints are revealed. Since evaluation prior to full constraint revelation cannot distinguish between α and α' , early commitment yields no informational advantage but strictly reduces reachable assignments. Hence early commitment is dominated. \square

This proposition formalizes a practical intuition: repair rewards patience. The presence of degradation and uncertainty forces agents to maintain slack and to treat materials, tools, and procedures as provisional. In this sense, degradation itself functions as a selection-pressure moderator. Physical systems impose constraints slowly and irreversibly, making premature evaluation maladaptive.

The same logic explains why maintenance cultures historically generate innovation. Maintenance is characterized by repeated encounters with small failures, incremental adjustments, and cumulative knowledge of system idiosyncrasies. Each intervention preserves or even expands the reservoir of function-neutral elements by decoupling components from their original design intent. Over time, these elements become candidates for exaptation in new contexts, often far removed from the original system.

By contrast, institutional environments that valorize pristine design and penalize visible failure suppress these dynamics. When degradation is treated as error rather than information, repair activities are marginalized or outsourced, and the exaptation window closes. The formal framework predicts that such environments will exhibit high initial efficiency but low long-term adaptability.

The central lesson is that repair is not merely a corrective activity but a generative one. It creates exactly the conditions required for creative intelligence: uncertain objectives, delayed evaluation, and reservoirs of partially decontextualized elements. Any theory of creativity that neglects repair and maintenance overlooks one of the most robust and historically validated sources of innovation.

13 Metric Saturation and the Collapse of the Exaptation Window

The previous sections established that exaptation requires delayed evaluation and that repair environments naturally provide such delays. We now turn to a complementary failure mode: the saturation of environments by metrics. Metric saturation refers to the condition in which evaluation is not merely early but continuous, frequent, and externally imposed, such that nearly every intermediate state is immediately judged according to a narrow performance signal. This section formalizes the claim that metric saturation collapses the exaptation window by converting exploratory trajectories into absorbing states.

Let evaluation be implemented not as a single terminal functional but as a stream of measurements applied at discrete intervals. In metric-saturated environments, the interval between evaluations is

small relative to the horizon required for repurposing. Examples include engagement metrics on platforms, continuous grading in education, and real-time performance dashboards in organizations.

Formally, let $\delta > 0$ denote the evaluation frequency, measured as the number of evaluation points per unit horizon. As $\delta \rightarrow \infty$, evaluation approaches a continuous-time limit. Each evaluation induces a projection of the current state and assignment onto an acceptability set determined by the metric. States outside this set are penalized or eliminated.

Definition 11 (Metric-Induced Absorbing States). Given an evaluation metric $M : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ and threshold θ , define the absorbing set

$$\mathcal{A}_{\text{abs}} := \{(s, \alpha) : M(s, \alpha) < \theta\}.$$

Once a trajectory enters \mathcal{A}_{abs} , it is either terminated or irreversibly redirected toward conforming states.

In practice, absorption may take the form of defunding, invisibility, demotion, or abandonment. The essential property is irreversibility: once absorbed, the trajectory can no longer explore repurposing paths.

We now state the central result.

Theorem 3 (Critical Evaluation Frequency). *Assume there exists a minimal time $\tau_{\text{ex}} > 0$ required to complete any exaptation transition from initial assignment α_0 to some $\alpha' \in \mathcal{G}$. If evaluation frequency satisfies $\delta > 1/\tau_{\text{ex}}$, then no exaptation transition can be completed without encountering evaluation. Moreover, if evaluation penalizes intermediate assignments required for exaptation, then the expected exaptation yield converges to zero as $\delta \rightarrow \infty$.*

Proof. If $\delta > 1/\tau_{\text{ex}}$, then at least one evaluation occurs in every interval of length τ_{ex} . Since exaptation requires traversing intermediate assignments over this interval, each such trajectory is subject to evaluation before completion. If the metric penalizes these intermediate assignments, the trajectory is absorbed before reaching α' . As evaluation frequency increases, the probability of surviving long enough to complete exaptation decreases monotonically and tends to zero in the continuous evaluation limit. \square

This theorem provides a mechanistic explanation for why environments saturated with metrics appear hostile to creativity. The problem is not that metrics are inaccurate or unfair, but that their temporal density is incompatible with the time required for repurposing. Even benign metrics become destructive when applied too frequently.

The result also explains why adding additional metrics often worsens outcomes. Multiple metrics increase the dimensionality of the absorbing set, enlarging \mathcal{A}_{abs} and accelerating absorption. Under such conditions, trajectories collapse rapidly into a small number of conforming patterns, producing the familiar homogenization observed on platforms and in institutions.

The framework also clarifies the relationship between metric saturation and Goodhart-like phenomena. When metrics are used as targets, agents optimize directly for the metric, bypassing

the underlying objective. In the present formalism, this corresponds to agents steering trajectories to remain within the acceptability set at each evaluation point, rather than exploring assignments that would yield higher long-term value. The collapse of exaptation is therefore not a secondary side effect of Goodharts Law but a primary structural consequence of metric frequency.

Finally, the analysis highlights a crucial asymmetry: removing metrics after saturation does not restore the exaptation window. Once reservoirs have collapsed and assignments have converged to absorbing states, the entropy required for repurposing is lost. This irreversibility explains why institutional reforms that relax evaluation after long periods of metric saturation often fail to recover creativity. The damage has already been done.

Metric saturation thus emerges as a central antagonist in the theory of creative intelligence. By accelerating evaluation beyond the timescale of exaptation, it converts open-ended exploratory systems into closed, brittle ones. Any attempt to preserve creativity at scale must therefore treat evaluation frequency, not merely evaluation accuracy, as a first-order design variable.

14 Temporal Asymmetry, Irreversibility, and the Loss of Option Value

A defining feature of exaptation-based creativity is its temporal asymmetry. While the preservation of function-neutral elements requires time and restraint, their destruction through premature evaluation is rapid and often irreversible. This section formalizes the claim that early selection induces losses in option value that cannot be recovered by later permissiveness, even if evaluative pressure is subsequently relaxed.

To make this precise, consider an assignment space \mathcal{A} representing possible function assignments of elements within a system. Let $\Omega_0 \subset \mathcal{A}$ denote the initial set of admissible assignments prior to evaluation. Evaluation induces a contraction operator $E : \mathcal{P}(\mathcal{A}) \rightarrow \mathcal{P}(\mathcal{A})$ such that $E(\Omega) \subseteq \Omega$, reflecting the elimination of assignments deemed unacceptable. Iterated evaluation produces a nested sequence

$$\Omega_0 \supseteq \Omega_1 \supseteq \Omega_2 \supseteq \dots,$$

where $\Omega_{t+1} = E(\Omega_t)$.

Exaptation depends on the existence of multiple viable assignments and the ability to traverse between them. A natural measure of this capacity is the entropy of the assignment set, $H(\Omega_t)$. Under mild assumptions on E , entropy is non-increasing over time. Once an assignment is removed, it cannot be recovered without external intervention that introduces new elements or relaxes constraints in a way that exceeds the original state.

Theorem 4 (Irreversibility of Early Evaluation). *If evaluation operator E is idempotent and strictly contractive on Ω_0 , then for any time $t > 0$, there exists no operator R acting solely on Ω_t such that $R(\Omega_t) = \Omega_0$. In particular, entropy loss induced by early evaluation is irreversible without external injection of new degrees of freedom.*

Proof. Since E is strictly contractive, there exists at least one assignment $\alpha \in \Omega_0$ such that $\alpha \notin \Omega_1$. By idempotence, $\alpha \notin \Omega_t$ for all $t \geq 1$. Any operator R acting only on Ω_t cannot generate elements outside Ω_t by definition. Hence α cannot be recovered, and Ω_0 is unreachable from Ω_t . Entropy loss follows immediately. \square

This result formalizes an intuition widely observed in practice: once creative latitude is eliminated, it cannot be reconstituted simply by removing constraints later. Educational systems that impose rigid curricula early and introduce flexibility only at advanced stages fail to recover lost generativity. Similarly, organizations that subject exploratory work to early performance review cannot compensate by offering later freedom; the reservoir has already been depleted.

Temporal asymmetry also explains why evaluation timing is more consequential than evaluation severity. A mild evaluative filter applied early can be more destructive than a stringent filter applied late, because early evaluation acts on a richer assignment space. By contrast, late evaluation operates on trajectories that have already consolidated viable exaptations, preserving their internal coherence even if many alternatives are subsequently rejected.

This asymmetry has direct implications for system design. If the goal is to preserve long-term creative capacity, evaluation must be deferred until after exaptation windows have closed naturally through interaction with real constraints. Artificially accelerating selection produces systems that appear efficient in the short term but are brittle under novelty.

The analysis also clarifies why appeals to resilience or adaptability often fail in highly optimized systems. Adaptation presupposes the existence of latent alternatives. Once early evaluation has collapsed these alternatives, adaptation becomes impossible without importing novelty from outside the system. In this sense, premature evaluation converts endogenous creativity into a dependence on exogenous shocks.

Temporal asymmetry thus emerges as a fundamental constraint on creative intelligence. Selection is easy to apply and hard to undo. Any architecture that ignores this asymmetry risks mistaking irreversible loss for temporary inefficiency. Preserving exaptation is therefore not merely a matter of allowing creativity at some point, but of respecting the temporal conditions under which creativity can exist at all.

15 Exaptation and Artificial Intelligence: Pretraining, Fine-Tuning, and Alignment Pressure

The framework developed thus far provides a precise lens through which to reinterpret recent advances and failures in artificial intelligence. Contemporary large-scale models are often described as exhibiting creativity, generality, or emergent intelligence. From the present perspective, these properties do not arise from superior optimization per se, but from an accidental preservation of exaptation buffers during training, followed by their partial destruction during deployment and alignment.

Modern machine learning pipelines are typically divided into two phases: pretraining and fine-tuning. Pretraining exposes a model to vast quantities of heterogeneous data under weak or diffuse objectives, while fine-tuning applies targeted optimization toward specific tasks or behaviors. Although this distinction is usually justified pragmatically, it admits a deeper interpretation in terms of selection pressure management.

During pretraining, evaluation is sparse and indirect. Loss functions operate at a statistical level, and no single datum is decisive. Representations are therefore shaped under conditions approximating function-neutrality: features are retained not because they solve a particular task, but because they co-occur across many contexts. This phase constructs a large reservoir of latent structure with high entropy and high option value.

Fine-tuning, by contrast, introduces dense, task-specific evaluation. Gradients become directional, loss landscapes sharpen, and representations are selectively pruned to satisfy narrow objectives. This transition corresponds exactly to the imposition of selection pressure analyzed earlier. Exaptation capacity is not created during fine-tuning; it is consumed.

We can formalize this intuition by modeling representation learning as the construction of a latent space \mathcal{Z} equipped with a family of decoders $\{\pi_T\}$ indexed by tasks T . Let $H(\mathcal{Z})$ denote the entropy of the latent space with respect to task-agnostic structure. Pretraining seeks to maximize mutual information between inputs and \mathcal{Z} subject to weak regularization, whereas fine-tuning minimizes loss for a specific T^* .

Proposition 5. *If fine-tuning minimizes task-specific loss without constraint on latent entropy, then $H(\mathcal{Z})$ is non-increasing and strictly decreases whenever task-irrelevant features are penalized. Consequently, the expected exaptation option value of \mathcal{Z} decreases monotonically with fine-tuning intensity.*

Proof. Task-specific loss assigns negative gradient to latent dimensions that do not contribute to performance on T^* . Absent entropy-preserving regularization, these dimensions are suppressed. Since exaptation option value depends on the availability of latent features for unforeseen tasks, reducing latent dimensionality strictly reduces expected option value. \square

This result explains a now-familiar empirical pattern: models often lose generality, robustness, or creativity when over-fine-tuned or aggressively aligned. The issue is not overfitting in the classical statistical sense, but premature collapse of the exaptation reservoir. Alignment procedures that apply dense human feedback at every step function analogously to metric saturation, accelerating evaluation beyond the timescale required for repurposing.

The framework also clarifies why large models appear more capable than smaller, more tightly optimized ones. Scale matters not because it enables deeper optimization, but because it increases the size of the function-neutral reservoir constructed during pretraining. A larger model can absorb more heterogeneity without immediate collapse, preserving exaptation potential across a wider range of contexts.

This interpretation has direct implications for AI governance. Efforts to make models safe, predictable, or aligned often proceed by increasing evaluation density and narrowing objectives.

While such measures may improve short-horizon control, they simultaneously reduce the systems capacity to respond intelligently to novel situations. The trade-off is structural, not accidental.

From the present perspective, the challenge of AI alignment is inseparable from the problem of selection pressure management. An aligned system that has lost its exaptation buffer may be safe but brittle; a system that retains exaptation capacity may be adaptable but unpredictable. Resolving this tension requires explicit architectural mechanisms that delay or compartmentalize evaluation, rather than applying it uniformly.

Artificial intelligence thus provides a concrete instantiation of the general theory developed in this essay. Creativity, adaptability, and intelligence emerge not from relentless optimization, but from sustained exposure to weakly evaluated environments followed by judicious, late-stage selection. Where evaluation is imposed too early or too frequently, intelligence collapses not because the system lacks power, but because its future has been foreclosed.

16 Evaluation Timing as Governance: Power, Control, and the Allocation of Futures

The preceding analysis has treated evaluation timing primarily as a technical variable affecting creative capacity. This section extends the argument by showing that evaluation timing is also a governance mechanism. Decisions about when evaluation occurs are not neutral implementation details; they allocate power by determining which futures are reachable and by whom. From this perspective, selection pressure management is inseparable from political and institutional structure.

To formalize this claim, consider a population of agents operating within a shared assignment space \mathcal{A} . Let each agent i have access to a subset $\mathcal{A}_i \subseteq \mathcal{A}$ determined by material resources, institutional permissions, and evaluative regimes. Let evaluation schedules be parameterized by a function $\tau_i(t)$ specifying the earliest time at which agent i 's actions are subject to binding evaluation.

An agent with a long evaluation delay τ_i can explore assignments over a larger region of \mathcal{A}_i before contraction occurs, while an agent with a short τ_i is forced into early convergence. The asymmetry in reachable assignments translates directly into an asymmetry in creative and strategic power.

Theorem 5 (Evaluation Timing and Reachable Futures). *Let $\mathcal{R}_i(\tau_i)$ denote the set of assignments reachable by agent i before evaluation-induced contraction. If $\tau_i > \tau_j$, then generically $\mathcal{R}_i(\tau_i) \supsetneq \mathcal{R}_j(\tau_j)$. Moreover, the difference in reachable sets grows superlinearly with the difference in evaluation delay under mild assumptions on assignment connectivity.*

Proof. Evaluation delay determines the maximum length of exploratory trajectories before contraction. In connected assignment spaces, reachable volume grows with trajectory length. Since contraction removes regions irreversibly, agents with shorter delays lose access to assignments that remain available to agents with longer delays. Superlinearity follows from branching in assignment paths. \square

This result reveals evaluation timing as a mechanism for concentrating creative capacity. Institutions that grant themselves long horizons while imposing short horizons on individuals effectively monopolize exaptation. Research laboratories, corporations, and states often reserve the right to experiment privately while demanding immediate accountability from workers, students, or users. The asymmetry is structural rather than conspiratorial.

The same mechanism explains why credentialing systems, while restrictive, historically functioned as horizon-extending institutions. By delaying evaluation through prolonged training and apprenticeship, they granted practitioners protected exaptation windows. However, when credentialing becomes purely gatekeeping without corresponding protection of exploratory time, it ceases to serve this function and instead reinforces hierarchy.

Digital platforms provide a particularly stark illustration. Platform operators enjoy long, opaque evaluation cycles insulated by scale and legal protections, while users are subjected to continuous evaluation via metrics, moderation, and algorithmic ranking. The platform thus occupies a region of \mathcal{A} inaccessible to its participants, enabling it to innovate and pivot while users converge into predictable patterns.

This asymmetry also clarifies why appeals to democratization often ring hollow. Expanding access to tools without expanding access to delayed evaluation merely increases the number of agents competing under the same collapsed horizon. Equality of opportunity in such environments is illusory, as the distribution of reachable futures remains skewed.

From a governance perspective, the critical question is therefore not who may speak or act, but under what evaluative schedule. Freedom without temporal protection is fragile; it permits expression but not exploration. Conversely, delayed evaluation without accountability risks unbounded drift. The challenge is to design institutions that distribute exaptation buffers without abandoning eventual selection.

The analysis suggests a reframing of institutional ethics. Rather than asking whether evaluation is fair or accurate, one must ask whether its timing preserves or forecloses future possibility. Control over evaluation schedules is control over the space of potential outcomes. In this sense, selection pressure management is a primary axis of power, and creative intelligence is inseparable from the politics of time.

17 Cultural Memory, Archives, and Long-Horizon Exaptation

The dynamics of exaptation and delayed evaluation extend beyond individual cognition and institutional design to the level of culture. Societies differ markedly in their capacity to preserve, recombine, and repurpose knowledge across generations. These differences can be analyzed within the present framework by treating cultural memory as a long-horizon reservoir of function-neutral elements subject to exceptionally delayed evaluation. Archives, libraries, notebooks, craft traditions, and informal repositories function as exaptation buffers operating on timescales far longer than those available to individuals or organizations.

Let \mathcal{K} denote a cultural knowledge space composed of artifacts, texts, techniques, and practices.

At any time t , a society maintains a subset $\mathcal{K}_t \subseteq \mathcal{K}$ that remains accessible. Evaluation operates by selectively preserving, canonizing, or discarding elements of \mathcal{K}_t based on prevailing criteria of relevance or utility. In cultures with strong archival traditions, the evaluation operator acts weakly and infrequently, allowing large regions of \mathcal{K}_t to persist without assigned function.

This persistence enables exaptation across historical discontinuities. Elements preserved for reasons unrelated to their future use may later become central when new constraints arise. Manuscripts copied for religious devotion become sources for scientific insight; artisanal techniques developed for local materials become templates for industrial processes; marginal theories become foundational under new empirical regimes. In each case, creative advance depends not on foresight but on preservation.

We can formalize this effect by modeling cultural memory as a multi-period reservoir subject to delayed contraction. Let E_T denote an evaluation operator applied at interval T , with T measured in generations rather than years. As T increases, the expected number of latent exaptation pathways grows.

Proposition 6. *Assume cultural knowledge elements have independent probabilities of future relevance under unknown environmental changes. Then the expected exaptation yield of a cultural reservoir is increasing in the evaluation interval T and in the size of the preserved knowledge set $|\mathcal{K}_t|$.*

Proof. Longer evaluation intervals allow more elements to persist through periods of apparent irrelevance. Since future relevance is independent of present evaluation criteria, preserving more elements increases the probability that at least one will match future constraints. Increasing T reduces the rate of premature elimination, raising expected yield. \square

This result explains why cultures that tolerate redundancy, apparent inefficiency, and archival excess often outperform more aggressively optimized societies over long horizons. What appears locally as waste or irrelevance functions globally as insurance against epistemic shock. Conversely, cultures that subject knowledge to frequent pruning based on short-term relevance systematically reduce their adaptive capacity.

The analysis also clarifies the fragility of cultural memory under modern conditions. Digitization, metric-driven curation, and algorithmic ranking introduce continuous evaluation into domains that historically operated under delayed selection. When archives are reorganized according to popularity, recency, or engagement, their function-neutrality collapses. Knowledge becomes visible only insofar as it performs, and exaptation potential is lost.

Importantly, the loss is again irreversible. Once materials are discarded, unindexed, or rendered inaccessible, later recognition of their value cannot recover the missing pathways without extraordinary effort. Cultural amnesia is therefore not merely a loss of information but a contraction of the future.

This perspective reframes preservation as an active cognitive strategy rather than a nostalgic impulse. To preserve artifacts, texts, and practices without knowing their future use is not to resist

progress but to enable it. Cultural memory functions as a collective exaptation buffer, distributing creative capacity across generations by refusing to collapse function prematurely.

In this light, the value of libraries, archives, and slow scholarship cannot be justified solely in terms of immediate utility. Their primary contribution lies in maintaining a space of latent possibility whose significance may only become apparent under future constraints. Societies that dismantle these buffers in the name of efficiency trade short-term clarity for long-term fragility.

Cultural exaptation thus completes the scale-free picture developed throughout this essay. From individual repair practices to institutional design and intergenerational memory, creative intelligence depends on the same underlying mechanism: the preservation of function-neutral elements under delayed evaluation. Where this mechanism is sustained, innovation persists. Where it is suppressed, the future narrows.

18 Rapid Apparent Evolution Under Shifting Selection Pressures

Microbial evolution is often described as “fast” or “accelerated,” particularly in bacteria and viruses subjected to rapidly changing environments such as immune systems, antibiotics, or antiviral drugs. This apparent speed is frequently attributed to short generation times or high mutation rates. While these factors are relevant, they are insufficient to explain the qualitative pattern observed in practice: microbes often appear to respond almost immediately to novel constraints, as though pre-adapted to conditions they have never previously encountered. The present framework provides a more precise explanation grounded in exaptation and selection pressure management.

Bacterial and viral populations maintain extraordinarily large reservoirs of genetic and phenotypic variation that are, at any given time, largely function-neutral. Many mutations are neutral or nearly neutral with respect to current environmental demands. Classical evolutionary theory treats this variation as background noise or genetic drift. From the perspective of exaptation, however, this variation constitutes a standing reservoir of latent function assignments awaiting future selection.

Let \mathcal{G} denote the space of possible genotypes and let \mathcal{P} denote the space of phenotypic effects. At any time t , a population occupies a subset $\mathcal{G}_t \subset \mathcal{G}$ whose members express phenotypes that are viable under existing constraints. Crucially, many elements of \mathcal{G}_t encode phenotypic capacities that are not currently expressed or evaluated. These capacities persist precisely because selection pressure is weak or absent with respect to them.

When environmental conditions change abruptly, the evaluation function changes faster than the population can generate new variation. Adaptation therefore proceeds not primarily through new mutation, but through the reassignment of function to pre-existing variants. What appears as rapid evolution is, in fact, rapid exaptation.

This process can be formalized as follows. Let $E_t : \mathcal{G} \rightarrow \{0, 1\}$ denote a selection operator encoding viability at time t . A sudden environmental change corresponds to a sharp change in E_t to E_{t+1} . If \mathcal{G}_t contains elements that were previously neutral under E_t but viable under E_{t+1} , then the population can shift rapidly without traversing mutational space.

Proposition 7. *If a population maintains a sufficiently large set of function-neutral genotypes under E_t , then under abrupt change to E_{t+1} the expected time to adaptation is bounded independently of mutation rate.*

Proof. Adaptation time depends on the existence of genotypes in \mathcal{G}_t that satisfy E_{t+1} . If such genotypes already exist, selection acts by amplification rather than exploration. Since amplification operates on existing variants, its timescale is governed by replication dynamics rather than mutation, yielding rapid apparent adaptation. \square

This result explains why bacteria and viruses often evade antibiotics or immune responses within a few generations. Resistance mechanisms such as efflux pumps, altered binding sites, or metabolic bypasses frequently predate the selective pressure that makes them advantageous. Their persistence prior to selection reflects delayed evaluation: in the absence of pressure, these traits are neither rewarded nor eliminated.

Importantly, environments characterized by fluctuating or heterogeneous selection pressures actively favor the preservation of such reservoirs. In microbial populations exposed to variable conditions, selection against latent capacities is weakened because future relevance is unpredictable. This produces populations with high exaptation potential. By contrast, stable environments encourage specialization and reduce latent diversity, making rapid adaptation less likely.

The same logic applies to viral quasispecies. RNA viruses, in particular, maintain clouds of closely related variants, many of which are suboptimal or deleterious under current conditions. These variants persist because selection is applied at the level of population viability rather than individual optimality. When the environment shifts, variants that were previously suppressed but not eliminated can become dominant almost instantaneously.

From this perspective, microbial “evolvability” is not a mysterious property but a consequence of selection pressure management imposed by ecological structure. Rapid apparent evolution emerges when large, weakly evaluated reservoirs are subjected to sudden constraint changes. The speed lies not in foresight, but in preserved optionality.

This analysis also clarifies why aggressive and continuous selection, such as high-dose antibiotic treatment, can paradoxically accelerate resistance. By imposing strong but narrow constraints, such regimes collapse some regions of genotypic space while leaving others untouched, favoring variants that escape the evaluated dimensions entirely. Selection pressure that is intense but myopic thus promotes exaptation along orthogonal axes.

The broader implication is that evolutionary speed is not solely a function of mutation rate or generation time, but of the temporal structure of evaluation. Systems that preserve large reservoirs of function-neutral variation under delayed or intermittent selection will appear extraordinarily adaptive when constraints shift. What is observed as rapid evolution is, in fact, the late-stage unveiling of possibilities that were already present.

This completes the biological grounding of the framework. The same mechanism that governs microbial adaptation under shifting environments also governs creativity, learning, institutional

failure, and cultural persistence. Across scales, intelligence emerges not from constant optimization, but from the protection of latent structure until selection can act meaningfully.

19 Antibiotic Resistance as Selection Pressure Mismanagement

Antibiotic resistance provides a concrete and policy-relevant case study of the general mechanism developed in this essay. Resistance is often framed as a problem of insufficient compliance, improper dosing, or microbial ingenuity. While these factors matter, they obscure a more fundamental issue: many antibiotic regimes impose selection pressures that are intense, narrow, and temporally misaligned with the dynamics of exaptation. As a result, they accelerate the emergence and fixation of resistant strains rather than suppressing them.

To situate the problem formally, consider a bacterial population occupying a genotypic reservoir \mathcal{G}_t with associated phenotypic effects relevant to survival under antibiotic exposure. Let E_A denote the evaluation operator induced by a particular antibiotic A , mapping genotypes to viability outcomes. Classical therapeutic logic assumes that sufficiently strong application of A will eliminate the population by collapsing \mathcal{G}_t to the empty set. In practice, however, E_A acts along a limited set of biochemical dimensions, leaving orthogonal dimensions unevaluated.

Resistance mechanisms frequently arise from these orthogonal dimensions. Traits such as efflux pumps, metabolic pathway redundancy, biofilm formation, and target modification often exist prior to antibiotic exposure and serve multiple functions unrelated to resistance. Under weak or absent selection pressure, these traits persist as function-neutral or mildly deleterious variants. Antibiotic application converts them into dominant survival strategies through exaptation.

This process can be formalized by decomposing genotypic effects into evaluated and unevaluated components. Let $\mathcal{G} = \mathcal{G}_A \times \mathcal{G}_\perp$, where \mathcal{G}_A represents traits directly targeted by antibiotic A , and \mathcal{G}_\perp represents traits orthogonal to its mechanism. Strong selection against \mathcal{G}_A rapidly collapses variation along that axis, but leaves \mathcal{G}_\perp largely untouched.

Theorem 6 (Resistance via Orthogonal Exaptation). *If a population contains variants in \mathcal{G}_\perp that confer survival under E_A , then increasing the strength of selection along \mathcal{G}_A strictly increases the relative fitness of those variants, accelerating resistance fixation.*

Proof. Let $g = (g_A, g_\perp)$ denote a genotype. Strong selection reduces the survival probability of all g_A not satisfying E_A , effectively projecting the population onto a narrow subset of \mathcal{G}_A . Variants differing only in g_\perp are unaffected by this projection. If some g_\perp confer survival under antibiotic exposure, their relative frequency increases monotonically as competitors are eliminated along \mathcal{G}_A . Hence stronger selection accelerates fixation. \square

This theorem explains why aggressive antibiotic use often backfires. By collapsing variation along the targeted axis while leaving latent resistance mechanisms intact, treatment regimes create ideal conditions for exaptation. What appears as microbial ingenuity is in fact a predictable outcome of selection pressure geometry.

Temporal structure further exacerbates the problem. Continuous high-dose treatment applies dense evaluation at every generation, eliminating any opportunity for intermediate states to persist. While this may reduce population size initially, it also ensures that only genotypes capable of surviving under constant pressure remain. Intermittent or heterogeneous exposure, by contrast, introduces temporal slack that can prevent the consistent amplification of any single resistance mechanism.

Importantly, resistance does not require that bacteria “anticipate” antibiotics. It emerges because evaluation is applied too narrowly and too early, collapsing function-neutral reservoirs in a way that favors orthogonal escape routes. From the present perspective, resistance is not a failure of microbial control but a failure of selection pressure management.

This analysis has direct implications for antibiotic policy. Strategies that vary selection pressures across time, space, or mechanism reduce the likelihood that any single exaptation pathway will dominate. Combination therapies, cycling protocols, and ecological approaches that preserve competitive diversity all function by disrupting the alignment between evaluation frequency and exaptation timescales.

More broadly, antibiotic resistance illustrates a general principle: when selection pressure is intense, continuous, and unidimensional, systems respond by exploiting latent degrees of freedom rather than by complying with intended constraints. Suppressing exaptation is neither feasible nor desirable; the only viable strategy is to design selection regimes that do not inadvertently privilege the very adaptations they seek to prevent.

Antibiotic resistance is therefore not an anomaly but a paradigmatic example of the framework developed throughout this essay. It demonstrates, in biological detail, how premature and narrowly focused evaluation collapses some possibilities while amplifying others, reshaping the future in ways that appear surprising only if the underlying mechanism is misunderstood.

20 Teleological Misreadings of Adaptation and the Illusion of Foresight

A persistent obstacle to understanding adaptation, creativity, and intelligence across domains is the tendency toward teleological misreading. When systems respond effectively to new constraints, observers frequently infer intention, foresight, or goal-directed planning. In biology, this manifests as the claim that organisms or populations “adapt in order to survive.” In institutions and technology, it appears as the belief that successful outcomes reflect superior strategy or insight. The framework developed in this essay shows that such interpretations are not merely imprecise but structurally misleading.

Teleological error arises when late-stage selection outcomes are projected backward onto the generative process that produced them. Exaptation makes this error particularly tempting, because the reassigned function often appears exquisitely matched to the new environment. Antibiotic resistance, for example, is commonly described as bacteria “figuring out” how to evade drugs. Yet

as shown in the previous section, resistance typically arises through the amplification of variants that pre-existed exposure and were not evaluated under prior conditions. The apparent intelligence lies not in anticipation, but in preserved optionality.

This misinterpretation can be formalized. Let α_0 denote an initial function assignment that is neutral under evaluation E_t , and let α_1 denote a reassigned function that becomes advantageous under E_{t+1} . Observers who only see α_1 under E_{t+1} may infer that the system optimized for E_{t+1} all along. In reality, α_0 persisted precisely because it was not optimized for any specific future evaluation. Teleology enters when persistence under non-evaluation is mistaken for preparation.

Proposition 8. *Given a system operating under delayed evaluation, any exapted outcome that survives late selection will appear teleologically optimized when observed without access to the prior evaluation schedule.*

Proof. Delayed evaluation ensures that intermediate assignments are not visible to selection-based filtering. Observers who condition only on surviving assignments under E_{t+1} sample from a distribution biased toward functional coherence. Without information about eliminated or neutral variants, Bayesian inference favors goal-directed explanations, even though the generative process was non-teleological. \square

This proposition explains why teleological narratives recur across evolutionary biology, cultural history, and institutional analysis. Survivorship bias combined with evaluation delay produces the illusion of intention. The stronger and more decisive the late-stage selection, the more compelling the illusion becomes.

Paul Feyerabends critique of methodological monism in *Against Method* is directly relevant here. Feyerabend argued that scientific progress does not proceed through orderly application of universal rules, but through historical processes rife with contradiction, redundancy, and apparent irrationality. Many ideas that later proved foundational survived only because they were protected from prevailing standards of evaluation. From the present perspective, Feyerabends epistemological anarchism can be reinterpreted as an argument for delayed evaluation as a precondition for intellectual exaptation.

Feyerabends insistence that theories must sometimes be defended against evidence is often misunderstood as an attack on rationality. In fact, it reflects a recognition of temporal asymmetry: early evidence is evaluated under existing conceptual frameworks, which are themselves subject to future revision. Premature evaluation collapses the space of possible theories before their latent affordances can be realized. What appears, in retrospect, as methodological deviance is in fact selection pressure management.

The teleological fallacy thus functions as a retrospective justification for optimization-centric narratives. By attributing success to foresight, institutions obscure the role of slack, waste, and protected deviation in producing that success. This misattribution then feeds back into policy, encouraging tighter control, earlier evaluation, and stricter adherence to methodall of which suppress the very processes that generated the outcome being celebrated.

In evolutionary discourse, this fallacy encourages the view that adaptation is driven by problem-solving intelligence rather than by population-level dynamics under delayed selection. In education,

it supports the belief that students should be optimized early for anticipated outcomes. In innovation policy, it legitimizes aggressive pruning of “unproductive” work. In each case, teleology serves as an ideological cover for premature selection.

Rejecting teleological explanations does not entail denying functionality or coherence. It requires recognizing that coherence is an emergent property of late-stage selection acting on reservoirs preserved under weak or absent evaluation. Exaptation produces results that look planned precisely because they were not.

This insight completes the conceptual arc of the essay. Across biological, cognitive, institutional, and cultural domains, apparent intelligence emerges from the same mechanism: the maintenance of latent structure under delayed evaluation, followed by decisive but temporally appropriate selection. Teleological narratives invert this order, mistaking outcomes for causes and optimization for creativity. Correcting this inversion is essential not only for theoretical clarity, but for the design of systems capable of sustaining intelligence over time.

21 Synthesis: Exaptation, Delayed Evaluation, and the Architecture of Creative Intelligence

This essay has developed a unified account of creative intelligence grounded in two inseparable components: exaptation as the generative operator and delayed evaluation as the enabling condition. Across biological, cognitive, institutional, and cultural scales, the same mechanism recurs. Systems capable of sustained intelligence preserve function-neutral reservoirs long enough for reassignment to occur and apply selection only after meaningful interaction with constraints. Where either component is absent, intelligence collapses into brittle optimization.

The analysis distinguished exaptation from transfer and generalization. Whereas transfer learning presupposes partial knowledge of future tasks and optimizes representations accordingly, exaptation assumes radical uncertainty and preserves optionality. This difference is mathematical rather than semantic: optimization for known distributions contracts representational entropy and irreversibly reduces option value for unknown futures, while creative intelligence depends on resisting such contraction until constraints are revealed.

Repair and maintenance were shown to be privileged sites of exaptation because they impose delayed evaluation by necessity. Failures reveal constraints incrementally, preventing premature commitment and forcing provisional use of materials, concepts, and procedures. In these environments, exaptation is routine rather than exceptional, whereas institutions that marginalize repair in favor of pristine design suppress historically robust sources of innovation.

The analysis of metric saturation demonstrated that evaluation frequency, not merely evaluation accuracy, determines creative viability. Dense and continuous evaluation causes exploratory trajectories to encounter absorbing states before repurposing can occur, producing homogenization and brittleness even under abundant resources. Crucially, the resulting loss of option value is temporally asymmetric and irreversible: later permissiveness cannot restore what early evaluation

has destroyed.

Artificial intelligence provided a contemporary instantiation of these dynamics. Pretraining functions as large-scale reservoir construction under weak evaluation, while fine-tuning and alignment impose late-stage selection. Models appear creative insofar as pretraining dominates and become brittle as evaluation density increases, revealing a structural tension between adaptability and control.

At the level of governance, evaluation timing was shown to allocate power by determining which futures are reachable. Agents or institutions granted long evaluation delays can explore broadly and exapt freely, while those subjected to early evaluation are forced into premature convergence. This asymmetry explains why access to tools alone does not democratize creativity: without temporal protection, expanded participation merely increases competition under collapsed horizons.

Cultural memory extended the framework across generations. Archives, libraries, and traditions function as long-horizon exaptation buffers, preserving materials without knowing their future use. Societies that tolerate redundancy and apparent inefficiency thereby maintain adaptive capacity under uncertainty, while those that subject knowledge to continuous relevance-based pruning systematically narrow their futures.

Biological evolution under shifting selection pressures provided decisive grounding. Rapid apparent evolution in bacteria and viruses arises not from foresight but from preserved reservoirs of function-neutral variation subjected to sudden constraint changes. Antibiotic resistance exemplified the consequences of selection pressure mismanagement, in which intense, narrow, and continuous evaluation amplifies orthogonal escape routes through exaptation.

Finally, the essay addressed teleological misreadings that project intention backward onto selection outcomes. Such narratives mistake late-stage coherence for foresight and obscure the role of slack, waste, and protected deviation. Feyerabend's critique of methodological rigidity was reinterpreted as an argument for delayed evaluation as a precondition for intellectual progress.

Taken together, these analyses converge on a single claim. Creative intelligence is not the product of relentless optimization or unbounded freedom, but of architectures that preserve function-neutral reservoirs under delayed evaluation. Selection is necessary, but its timing is decisive: when evaluation precedes exploration, futures collapse; when it follows exploration, coherence emerges.

The implications are both theoretical and practical. Systems intended to sustain intelligence under uncertainty must explicitly manage selection pressure by resisting early metricization, preserving slack, tolerating apparent waste, and distributing exaptation buffers rather than concentrating them. Exaptation and delayed evaluation are not optional enhancements to creativity; they are its necessary conditions, and their preservation determines whether a system retains a future at all.

A Formal Definitions

Let \mathcal{A} denote an assignment space of function mappings. Let $\Omega \subseteq \mathcal{A}$ denote the set of admissible assignments. Let $E : \mathcal{P}(\mathcal{A}) \rightarrow \mathcal{P}(\mathcal{A})$ denote an evaluation operator.

Definition 12 (Function-Neutral Reservoir). A reservoir Ω is function-neutral if for all $\alpha \in \Omega$, there exists no unique evaluation E such that α is optimal with respect to E .

Definition 13 (Exaptation). An exaptation is a mapping $\phi : \alpha \mapsto \alpha'$ such that α' performs a function not included in the design intent of α .

A.1 Entropy and Option Value

Let $H(\Omega)$ denote the Shannon entropy of assignments in Ω .

Definition 14 (Option Value). The option value of a reservoir Ω is defined as

$$V(\Omega) = \mathbb{E}_E \left[\max_{\alpha \in \Omega} U_E(\alpha) \right]$$

where U_E is utility under evaluation E .

Lemma 1. *If $\Omega_1 \subset \Omega_0$, then $V(\Omega_1) \leq V(\Omega_0)$.*

Proof. Immediate from set inclusion. □

A.2 Evaluation-Induced Contraction

Let $\Omega_{t+1} = E(\Omega_t)$.

Theorem 7. *If E is strictly contractive, then*

$$H(\Omega_{t+1}) < H(\Omega_t).$$

A.3 Irreversibility

Theorem 8. *If E is idempotent and contractive, no operator R exists such that*

$$R(E(\Omega)) = \Omega$$

unless R introduces elements not in $E(\Omega)$.

A.4 Exaptation Time Constraint

Let τ_{ex} denote minimal exaptation duration. Let δ denote evaluation frequency.

Theorem 9. *If $\delta > 1/\tau_{\text{ex}}$, exaptation probability tends to zero.*

A.5 Metric Saturation

Define absorbing set

$$\mathcal{A}_{\text{abs}} = \{\alpha \in \Omega : M(\alpha) < \theta\}.$$

A.6 Evolutionary Reservoirs

Let \mathcal{G} denote genotype space. Let $E_t : \mathcal{G} \rightarrow \{0, 1\}$ denote viability.

Theorem 10. *If $\exists g \in \mathcal{G}_t$ such that $E_{t+1}(g) = 1$ and $E_t(g) = 0$, adaptation time is bounded independently of mutation rate.*

A.7 Antibiotic Resistance

Decompose $\mathcal{G} = \mathcal{G}_A \times \mathcal{G}_\perp$.

Theorem 11. *Selection on \mathcal{G}_A increases fixation probability of advantageous \mathcal{G}_\perp variants.*

A.8 Governance and Horizon Allocation

Let τ_i denote evaluation delay for agent i .

Theorem 12. *If $\tau_i > \tau_j$, then generically*

$$|\mathcal{R}_i| > |\mathcal{R}_j|.$$

A.9 Teleological Illusion

Theorem 13. *Late-stage selection produces apparent optimization under incomplete observation.*

A.10 Summary Identity

$$\text{Creative Intelligence} = \text{Exaptation} + \text{Delayed Evaluation}.$$

References

- [1] Chis-Ciure, R., and Levin, M. (2025). Cognition all the way down 2.0: neuroscience beyond neurons in the diverse intelligence era. *Synthese*, 206, 257. <https://doi.org/10.1007/s11229-025-05319-6>
- [2] Feyerabend, P. (1975). *Against Method: Outline of an Anarchistic Theory of Knowledge*. Verso, London.
- [3] Zuboff, S. (2019). *The Age of Surveillance Capitalism*. PublicAffairs, New York.
- [4] Doctorow, C. (2023). The enshittification of TikTok. *Pluralistic*. Online essay.
- [5] Arendt, H. (1963). *Eichmann in Jerusalem: A Report on the Banality of Evil*. Viking Press, New York.
- [6] Schüll, N. D. (2012). *Addiction by Design: Machine Gambling in Las Vegas*. Princeton University Press, Princeton.
- [7] Gould, S. J., and Vrba, E. S. (1982). Exaptation: A missing term in the science of form. *Paleobiology*, 8(1), 4–15.
- [8] Mayr, E. (1982). *The Growth of Biological Thought*. Harvard University Press, Cambridge.
- [9] Goodhart, C. A. E. (1975). Problems of monetary management: The U.K. experience. *Papers in Monetary Economics*, Reserve Bank of Australia.
- [10] Polanyi, M. (1966). *The Tacit Dimension*. University of Chicago Press, Chicago.
- [11] Zupančič, A. (2024). *Disavowal*. Polity Press, Cambridge.
- [12] Holland, J. H. (1992). *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge, MA.
- [13] Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467–482.
- [14] Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science*, 43(7), 934–950.
- [15] Kauffman, S. A. (1993). *The Origins of Order*. Oxford University Press, Oxford.