

The Threshold of Deferral

A Crossover Theorem for Maintenance and Repair

Flyxion

2026

Abstract

This essay distinguishes continuous maintenance from discontinuous repair and derives a crossover condition governing when deferred correction becomes more costly than distributed intervention. Under compounding deviation dynamics and convex correction costs, maintenance and repair occupy distinct operational regimes separated by a threshold of deferral. The central result is a separation between two questions usually run together. Whether a deferral window opens *immediately* is governed entirely by the ratio of fixed overheads — continuous monitoring cost against per-event repair-initiation cost — independent of how fast deviation compounds or how convex correction cost is. Whether a deferral window can open *later*, under the opposite overhead ratio, is governed instead by a second, purely geometric comparison against a critical repair-overhead value built from the location of the cost differential's own unconstrained minimum. Together these two comparisons yield a complete, exhaustive classification of deferral regimes, including the fact — not obvious in advance — that late-opening deferral windows are finite intervals rather than permanent cutovers. This essay is a short companion to *Recursive Continuation*, refining the repair capacity term R_t of that essay's admissibility apparatus into its continuous (maintenance) and discontinuous (repair) constituents, and is intended to be read alongside it.

1 Two Kinds of Correction

Recursive Continuation treats repair capacity R_t as a single scalar offset against a convex maintenance burden $C(K_t)$, sufficient to establish that sustained viability requires $R_t \geq C(K_t)$ and that, past a threshold K^* , marginal repair investment dominates marginal capability investment (the Repair Dominance Theorem). That treatment was adequate for its purpose but conflated two operationally distinct mechanisms under one variable.

Repair, in the sense of that essay's Corollary 12.2, is conditional: it activates only once a candidate state has already left the viability manifold A , and its function is to return the trajectory to A after the fact. Maintenance is a different object. It acts on every step regardless of whether the current state has left A , and its function is to keep deviation small enough that repair is never triggered at all. The two are not different quantities of the same thing; they are different *kinds* of intervention, distinguished by their triggering condition rather than by their size.

Definition 1.1 (Decomposition of repair capacity). *Let R_t be the repair capacity of Recursive Continuation, Definition 10.3. Write*

$$R_t = M_t + \rho_t,$$

where M_t is maintenance, applied unconditionally at every step to reduce accumulated deviation before it compounds, and ρ_t is repair in the sense of Corollary 12.2 of that essay, applied only when $x_t \notin A$.

Remark. *This decomposition inherits the triggering structure of Corollary 12.2 without modification: ρ_t is precisely the map ρ of that corollary, restricted to the case where it is actually invoked. What this essay adds is a dynamical model of the deviation ρ acts on, and a cost comparison between paying M_t continuously and paying ρ_t occasionally.*

The question this essay answers is not whether maintenance is preferable to repair in general — that question, posed without further structure, collapses into an immediate consequence of convex cost and is not worth an essay. The question is under what conditions a rational system would choose one over the other, and where the boundary between those conditions lies. The answer, stated in full in §3, is a separation into two thresholds governed by disjoint mechanisms: *whether* an immediate deferral regime exists is governed entirely by fixed overheads, while *whether* a later one can open instead is governed by a second, purely geometric comparison. Those are different questions with different answers, and the bulk of this essay's technical content is owed to keeping them apart and resolving each in full.

2 Compounding Deviation

Definition 2.1 (Compounding deviation dynamics). *Let $d_t \geq 0$ denote accumulated deviation from an admissible reference state at time t . Let $a > 0$ be the deviation introduced at each step absent correction, and let $m_t \geq 0$ be maintenance applied before compounding occurs. Deviation evolves as*

$$d_{t+1} = (1 + \delta)(d_t - m_t) + a, \quad \delta \geq 0,$$

where δ is the delay-amplification rate: the rate at which uncorrected deviation compounds between steps.

Two limiting regimes are worth isolating before proceeding. If $m_t = d_t$ at every step (deviation fully corrected each period), $d_{t+1} = a$ for all t : deviation never accumulates, and the system pays a

constant per-step maintenance cost. If $m_t = 0$ for T consecutive steps and correction is deferred entirely, deviation accumulates according to a geometric sum,

$$D_T = a \sum_{k=0}^{T-1} (1 + \delta)^k = a \frac{(1 + \delta)^T - 1}{\delta}, \quad \delta > 0,$$

with $D_T = aT$ recovered in the limit $\delta \rightarrow 0$.

Remark. The case $\delta = 0$ is not a degenerate special case to be dismissed; it is the regime in which deviation merely accumulates without active amplification (a backlog of unrelated small defects, for instance), and it is the regime treated by §4's baseline before compounding is reintroduced. The essay's substantive content is what changes once $\delta > 0$ is allowed.

3 Cost Functionals and the Existence Proposition

Definition 3.1 (Maintenance and repair cost). Let correction cost be convex in the size of the deviation corrected, $c(d) = \alpha d^p$ for $\alpha > 0$, $p > 1$. Let $\kappa_M \geq 0$ be the fixed overhead of maintaining continuous correction capacity (monitoring, inspection, standing infrastructure) and $\kappa_R \geq 0$ the fixed overhead of initiating a single discontinuous repair event, independent of the size of the deviation repaired. Over a horizon of T steps, define

$$C_M(T) = T(\kappa_M + \alpha a^p), \quad C_R(T) = \kappa_R + \alpha \left(\frac{a}{\delta}\right)^p ((1 + \delta)^T - 1)^p, \quad \delta > 0.$$

$C_M(T)$ corrects the fixed per-step deviation a at every step, paying the maintenance overhead κ_M each time. $C_R(T)$ defers all correction to a single event at the end of the horizon, paying the repair overhead κ_R once and correcting the full compounded deviation D_T in one convex-cost operation.

Lemma 3.2 (Convexity of the cost differential). Let $D(T) = C_R(T) - C_M(T)$. Then D is convex on $T \geq 0$.

Proof. $g(T) = (1 + \delta)^T - 1$ satisfies $g''(T) = [\ln(1 + \delta)]^2 (1 + \delta)^T > 0$, so g is convex increasing. The map $u \mapsto u^p$ is convex increasing for $p > 1$, $u \geq 0$, and the composition of a convex increasing function with a convex function is convex; hence $f(T) = g(T)^p$ is convex, and $C_R(T) = \kappa_R + \alpha (a/\delta)^p f(T)$ is convex. $C_M(T)$ is affine. The difference of a convex function and an affine function is convex. \square

Convexity of D does not by itself guarantee a single, economically meaningful crossing. A convex function with $D(0) > 0$ and $D(T) \rightarrow \infty$ may have zero, one, or two roots depending on whether its minimum falls below zero, and if it does fall below zero, the two roots straddle the minimum. The following observation resolves this cleanly for the domain that matters.

Since D is convex, its derivative D' is monotone non-decreasing, and D therefore has a unique (possibly improper) global minimizer $T_{\min} \in [0, \infty]$, decreasing before T_{\min} and increasing after. Differentiating C_R and evaluating at $T = 1$ gives

$$D'(1) = \alpha a^p \left[\frac{p(1 + \delta) \ln(1 + \delta)}{\delta} - 1 \right] - \kappa_M,$$

so $D'(1) \geq 0$ — equivalently $T_{\min} \leq 1$ — holds exactly when

$$\kappa_M \leq \alpha a^p \left[\frac{p(1 + \delta) \ln(1 + \delta)}{\delta} - 1 \right]. \quad (\text{ER})$$

Definition 3.3 (Early Recovery Condition). *A system satisfies the Early Recovery Condition (ER) if $T_{\min} \leq 1$, equivalently if inequality (ER) holds. Under (ER), the unconstrained minimum of D occurs at or before the first unit of horizon, so that D is already on its ascending branch throughout the region $T \geq 1$ where threshold existence is decided: the system has already passed the point of maximum relative advantage for deferred repair by the time the shortest horizon under consideration is reached.*

Proposition 3.4 (Existence and Uniqueness of the Threshold of Deferral). *Evaluate D at $T = 1$: since $((1 + \delta)^1 - 1)^p = \delta^p$, the a, α, δ, p -dependence in $C_R(1)$ collapses exactly, giving*

$$D(1) = \kappa_R - \kappa_M.$$

Then:

- (i) *If $\kappa_R < \kappa_M$, there exists a unique $T^* > 1$ such that $C_R(T) < C_M(T)$ for $1 \leq T < T^*$ and $C_R(T) > C_M(T)$ for $T > T^*$. This holds unconditionally, regardless of δ, p, a, α .*
- (ii) *If $\kappa_R \geq \kappa_M$ and the Early Recovery Condition (ER) holds, no threshold exists in $T \geq 1$: $C_R(T) \geq C_M(T)$ for all $T \geq 1$, and maintenance weakly dominates repair from the shortest horizon onward.*

Proof. By the Lemma, D is convex, with $D(0) = \kappa_R > 0$ (taking $C_M(0) = 0$) and $D(T) \rightarrow \infty$ as $T \rightarrow \infty$, since C_R grows at rate $(1 + \delta)^{pT}$ against C_M 's linear growth. Convexity of D implies D has at most two zeros on $[0, \infty)$: D' is monotone, so D is strictly decreasing then strictly increasing (or monotone throughout), and a function of this shape crosses any fixed level at most twice.

(i) Suppose $\kappa_R < \kappa_M$, so $D(1) < 0$. Since $D(0) > 0 > D(1)$, one zero of D lies in $(0, 1)$, on the descending branch. Since $D(T) \rightarrow \infty$ and D has at most two zeros total, the second zero — if D does not remain negative forever, which it cannot since $D(T) \rightarrow \infty$ — lies on the ascending branch, necessarily after T_{\min} , and since $D(1) < 0$ it lies strictly after $T = 1$. Call it $T^* > 1$: $D(T) < 0$ for $1 \leq T < T^*$ and $D(T) > 0$ for $T > T^*$, giving exactly the stated dichotomy. No further zero is possible, so T^* is unique. This conclusion required only $D(0) > 0$, $D(1) < 0$, and convexity — not the Early Recovery Condition.

(ii) Suppose $\kappa_R \geq \kappa_M$, so $D(1) \geq 0$, and suppose in addition that the Early Recovery Condition holds, i.e. $T_{\min} \leq 1$. Because D is decreasing on $[0, T_{\min}]$ and increasing on $[T_{\min}, \infty)$, and $T_{\min} \leq 1$, D is non-decreasing throughout $[1, \infty)$. Since $D(1) \geq 0$, monotonicity gives $D(T) \geq D(1) \geq 0$ for all $T \geq 1$, i.e. $C_R(T) \geq C_M(T)$ throughout the domain of interest. \square

Remark (Why (ii) needs a hypothesis that (i) does not). *The asymmetry between the two parts is real, not a proof artifact. Part (i) needs only that $D(1) < 0$ together with convexity to locate a single ascending zero after $T = 1$; it does not need to know where T_{\min} falls, because $D(1) < 0$ already certifies that $T = 1$ lies at or after the descending branch has bottomed out below zero, and everything past a negative dip must eventually cross back up exactly once. Part (ii) is not the mirror image of this argument: $D(1) \geq 0$ alone does not rule out D dipping below zero after $T = 1$, which happens precisely when $T_{\min} > 1$ — that is, when compounding is aggressive enough that D 's unconstrained minimum has not yet been reached by $T = 1$. A numerical instance: $\delta = 2, p = 2, a = \alpha = 1, \kappa_M = \kappa_R = 5$ satisfies $\kappa_R \geq \kappa_M$ (with equality, so $D(1) = 0$) but violates the Early Recovery Condition, and in this instance $D(1.1) < 0$: a genuine deferral window opens shortly after $T = 1$ despite $\kappa_R \geq \kappa_M$. The Early Recovery Condition is exactly what excludes this.*

Remark (Scope of Proposition 3.4). *Proposition 3.4 establishes the paper's central separation result — existence of a deferral window opening exactly at $T = 1$ is governed entirely by κ_R versus κ_M , independent of*

δ, p, a, α — as an unconditional fact in one direction (i) and as a conditional fact in the other (ii), conditional on the Early Recovery Condition. What remains is the case $\kappa_R \geq \kappa_M$ together with failure of (ER): the late-recovery regime, illustrated by the numerical instance above, in which $D(1) \geq 0$ but D can still dip below zero later, since its unconstrained minimum has not yet been reached by $T = 1$. The next result resolves this regime completely.

3.1 The Late-Recovery Regime: A Completion Theorem

The late-recovery question reduces entirely to the sign of D at its own unconstrained minimum, $D(T_{\min})$. If $D(T_{\min}) \geq 0$, D never goes negative on $T \geq 1$ at all, since T_{\min} is where D is smallest; if $D(T_{\min}) < 0$, convexity guarantees D crosses zero exactly twice, straddling T_{\min} , opening a genuine — and, as shown below, *finite* — deferral window. Locating T_{\min} and evaluating D there in closed form completes the classification left open above.

Lemma 3.5 (Location of the Unconstrained Minimum). *For any $\kappa_M, \delta, p, a, \alpha > 0$, there is a unique $x^* > 0$ solving*

$$\alpha \left(\frac{a}{\delta}\right)^p p \ln(1 + \delta) x^{p-1} (1 + x) = \kappa_M + \alpha a^p, \quad (\dagger)$$

and $T_{\min} = \ln(1 + x^*) / \ln(1 + \delta)$ is the unique global minimizer of D .

Proof. Substituting $x = (1 + \delta)^T - 1$ and differentiating D gives $D'(T) = 0$ iff (\dagger) holds with $x = x(T)$. The left side of (\dagger) equals $\alpha(a/\delta)^p p \ln(1 + \delta)[x^{p-1} + x^p]$, a sum of two strictly increasing functions of x on $[0, \infty)$ (since $p - 1 > 0$ and $p > 0$), hence itself strictly increasing from 0 to ∞ ; a unique $x^* > 0$ therefore solves (\dagger) for any positive right-hand side. Since D is convex (Lemma 3.2), its unique stationary point is its global minimizer. \square

Lemma 3.6 (Positivity of the Bracket Term). *For $p > 1$ and $x > 0$, $p(1 + x) \ln(1 + x) > x$.*

Proof. Let $\psi(x) = x - p(1 + x) \ln(1 + x)$. Then $\psi(0) = 0$ and $\psi'(x) = 1 - p \ln(1 + x) - p \leq 1 - p < 0$ for all $x \geq 0$, since $\ln(1 + x) \geq 0$ and $p > 1$. Hence ψ is strictly decreasing from $\psi(0) = 0$, so $\psi(x) < 0$ for all $x > 0$, i.e. $x < p(1 + x) \ln(1 + x)$. \square

Proposition 3.7 (Late-Recovery Completion Theorem). *Let x^* be given by Lemma 3.5, and define*

$$\kappa_R^{\text{crit}} := \alpha \left(\frac{a}{\delta}\right)^p (x^*)^{p-1} \left[p(1 + x^*) \ln(1 + x^*) - x^* \right].$$

Then $\kappa_R^{\text{crit}} > 0$, and

$$D(T_{\min}) = \kappa_R - \kappa_R^{\text{crit}}.$$

Consequently, in the late-recovery regime ($\kappa_R \geq \kappa_M$, $x^* > \delta$):

- (i) If $\kappa_R < \kappa_R^{\text{crit}}$, there exist $1 < T_1^* < T_{\min} < T_2^*$ such that $C_R(T) < C_M(T)$ for $T \in (T_1^*, T_2^*)$ and $C_R(T) \geq C_M(T)$ otherwise. Deferral is rational only within this finite window, not for all sufficiently large T .
- (ii) If $\kappa_R \geq \kappa_R^{\text{crit}}$, no such window exists: $C_R(T) \geq C_M(T)$ for all $T \geq 1$.
- (iii) If $\kappa_R = \kappa_R^{\text{crit}}$, D touches zero at T_{\min} without crossing (tangency).

Proof. Positivity of κ_R^{crit} is immediate from Lemma 3.6. For the identity, write $A = \kappa_M + \alpha a^p$; then

$$D(T_{\min}) = \kappa_R + \alpha \left(\frac{a}{\delta}\right)^p (x^*)^p - A T_{\min}.$$

Using $T_{\min} \ln(1 + \delta) = \ln(1 + x^*)$ and solving (†) for A ,

$$\begin{aligned} A T_{\min} &= A \frac{\ln(1 + x^*)}{\ln(1 + \delta)} \\ &= \alpha \left(\frac{a}{\delta}\right)^p p \ln(1 + \delta) (x^*)^{p-1} (1 + x^*) \cdot \frac{\ln(1 + x^*)}{\ln(1 + \delta)} \\ &= \alpha \left(\frac{a}{\delta}\right)^p p (x^*)^{p-1} (1 + x^*) \ln(1 + x^*), \end{aligned}$$

the factor $\ln(1 + \delta)$ cancelling exactly. Substituting back,

$$D(T_{\min}) = \kappa_R + \alpha \left(\frac{a}{\delta}\right)^p (x^*)^{p-1} \left[x^* - p(1 + x^*) \ln(1 + x^*) \right] = \kappa_R - \kappa_R^{\text{crit}}.$$

For (i): $D(T_{\min}) < 0$ together with $D(1) \geq 0$ (since $\kappa_R \geq \kappa_M$) and convexity gives one zero $T_1^* \in (1, T_{\min})$ on the descending branch and, since $D(T) \rightarrow \infty$, a second zero $T_2^* \in (T_{\min}, \infty)$ on the ascending branch; $D < 0$ strictly between them and $D \geq 0$ outside, by convexity's single-dip property. For (ii): $D(T_{\min}) \geq 0$ is the global minimum value of D , so $D(T) \geq 0$ everywhere. (iii) is the boundary case of the same identity. \square

Remark (Consistency at the Early Recovery boundary). *At $x^* = \delta$ (the boundary of the Early Recovery Condition, $T_{\min} = 1$), (†) reduces exactly to $\kappa_M = \alpha a^p \left[\frac{p(1+\delta)\ln(1+\delta)}{\delta} - 1 \right]$ — the equality case of (ER) — and substituting $x^* = \delta$ into κ_R^{crit} gives, after using this identity to simplify, $\kappa_R^{\text{crit}} = \kappa_M$ exactly. The Late-Recovery Completion Theorem therefore glues continuously onto Proposition 3.4 at the boundary where the two results meet, rather than introducing a second, incompatible notion of threshold.*

Corollary 3.8 (Full Classification of Deferral Regimes). *With x^* and κ_R^{crit} as above:*

Regime	Criterion	Deferral-favorable set
Immediate threshold	$\kappa_R < \kappa_M$	$[1, T^*)$
Early-recovery, no threshold	$\kappa_R \geq \kappa_M, x^* \leq \delta$	\emptyset
Late-recovery, window	$\kappa_R \geq \kappa_M, x^* > \delta, \kappa_R < \kappa_R^{\text{crit}}$	(T_1^*, T_2^*)
Late-recovery, no window	$\kappa_R \geq \kappa_M, x^* > \delta, \kappa_R \geq \kappa_R^{\text{crit}}$	\emptyset
Critical surface	$\kappa_R = \kappa_R^{\text{crit}}$ (late) or $\kappa_R = \kappa_M$ (immediate)	tangency

(Immediate threshold: half-line cutoff. Late-recovery window: finite interval, as established in Proposition 3.7.) This classification is exhaustive and exact given the model of §2–3; no case remains open.

Remark (Two thresholds, two kinds of governance). *The classification separates cleanly into two comparisons governed by disjoint mechanisms. Whether deferral is attractive immediately is decided by κ_R versus κ_M — a comparison of fixed overheads, fixed by institutional or organizational structure, independent of compounding or curvature. Whether deferral becomes attractive later, having been unattractive at $T = 1$, is decided by κ_R versus κ_R^{crit} — a comparison against a quantity built entirely from where the compounding curve's own minimum sits, and hence governed by δ, p, a , and α rather than by overhead structure directly.*

The first threshold is an overhead phenomenon; the second is a geometric one. A further asymmetry worth noting explicitly: the immediate-threshold regime yields a permanent cutover (maintenance dominates for good once $T > T^*$), whereas the late-recovery regime yields only a bounded window of favorable deferral, closing again as T grows past T_2^* — deferral that becomes attractive late does not, on this model, stay attractive indefinitely.

4 Locating the Threshold: An Asymptotic Estimate

Proposition 3.4 establishes that T^* exists when $\kappa_R < \kappa_M$ but gives no closed form for its value, since $C_R(T) = C_M(T)$ is transcendental in T once $\delta > 0$. This section derives a perturbative estimate, proceeding in two stages: first the $\delta = 0$ baseline, then a first-order correction restoring compounding.

4.1 The $\delta \rightarrow 0$ Baseline

Write $A = \kappa_M + \alpha a^p$ and $B = \alpha a^p$. As $\delta \rightarrow 0$, $C_R(T) \rightarrow \kappa_R + BT^p$, and the crossover equation $AT = \kappa_R + BT^p$ is exactly solvable at $\kappa_R = 0$:

$$T_1 = \left(\frac{A}{B}\right)^{1/(p-1)} = \left(1 + \frac{\kappa_M}{\alpha a^p}\right)^{1/(p-1)}.$$

Remark. T_1 depends only on the ratio of fixed maintenance overhead to per-step correction cost and on the curvature p . Larger p compresses T_1 toward 1: strongly convex correction cost punishes any delay almost immediately, independent of how large κ_M is in absolute terms.

4.2 First-Order Correction in κ_R

Setting $T^* = T_1 + \Delta T$ in $AT = \kappa_R + BT^p$, linearizing, and using $AT_1 = BT_1^p$ to cancel the zeroth-order terms gives

$$\Delta T = -\frac{\kappa_R}{A(p-1)}.$$

Fixed repair overhead therefore pulls the threshold down from the $\kappa_R = 0$ baseline, as expected: cheap-to-trigger repair narrows the range over which deferral remains attractive.

4.3 First-Order Correction in δ

Using $\frac{(1+\delta)^T - 1}{\delta} \approx T + \frac{T(T-1)}{2}\delta$ for small δ , and perturbing $T = T_0 + \delta T'$ around the κ_R -corrected baseline T_0 , linearizing in δ with $BT_0^{p-1} \approx A$ to leading order gives

$$T' \approx -\frac{pT_1(T_1 - 1)}{2(p-1)}.$$

4.4 Combined Estimate

Proposition 4.1 (Asymptotic Location of the Threshold). *For $\kappa_R \ll A(p-1)(T_1 - 1)$ and $\delta T_1 \ll 1$,*

$$T^* \approx T_1 - \frac{\kappa_R}{(p-1)(\kappa_M + \alpha a^p)} - \delta \cdot \frac{pT_1(T_1 - 1)}{2(p-1)}, \quad T_1 = \left(1 + \frac{\kappa_M}{\alpha a^p}\right)^{1/(p-1)}.$$

Remark (Reading the estimate). *All three terms move the threshold in an interpretable direction, and the interaction between the second and third is the estimate’s most useful qualitative content. The δ -correction scales with $T_1(T_1 - 1)$: a system with a large pure-curvature baseline T_1 — one that would, on curvature alone, tolerate a long deferral window — loses threshold length faster under compounding than a system with a small baseline. Deferral tolerance is not eroded uniformly by compounding; it is eroded fastest in precisely the systems that most appear, on a static reading, to be able to afford it.*

Remark (Domain of validity). *As $\kappa_R \rightarrow \kappa_M^-$, Proposition 3.4 guarantees $T^* \rightarrow 1^+$ exactly, but the asymptotic estimate above is not valid in this limit, since it was derived by perturbing around T_1 , not around $T = 1$. Near this boundary the threshold should be read off the exact implicit equation $C_M(T) = C_R(T)$, numerically if necessary, rather than from the perturbative formula.*

5 Four Illustrations

The examples below are chosen to isolate one governing parameter each, following the same structure as the K^* illustrations of *Recursive Continuation*, §11.1: a worked case for δ , one for the existence condition of Proposition 3.4, one for p , and a capstone in which all parameters interact.

5.1 Biological Homeostasis: Compounding and δ

An untreated infection, an unmanaged vascular condition, or an unaddressed metabolic imbalance is characterized by active endogenous amplification: deviation does not merely persist, it grows, because the pathological process itself alters the conditions under which it continues to progress. This is a large- δ regime in the model’s terms.

The relevant point is not that such conditions are untreatable once advanced, but that the deferral window collapses rapidly, and does so fastest for exactly the class of conditions that would appear, on a static assessment, most tolerable to postpone. A slow-growing condition with a large notional T_1 is precisely the case the asymptotic estimate identifies as most sensitive to δ : the same growth rate that seems negligible at diagnosis erodes a seemingly generous threshold disproportionately. A system that reasons about deferral using only the curvature of eventual treatment cost, without separately accounting for δ , will systematically overestimate how much time it has.

5.2 Software Maintenance: The Existence Condition

A dormant, non-propagating software defect illustrates not the location of T^* but whether a deferral window exists at all. Such a defect typically has $\delta \approx 0$: it does not actively worsen while unaddressed. What determines whether deferring it is rational is the comparison between the fixed overhead of continuous correction capacity — CI infrastructure, scheduled code review, standing monitoring, all maintained whether or not anything is currently wrong, contributing to κ_M — against the fixed overhead of triggering a single repair event, typically small: opening a ticket, isolating a commit, shipping a hotfix, contributing to κ_R .

For an isolated defect with negligible compounding, $\kappa_R < \kappa_M$ is routinely satisfied, and Proposition 3.4 guarantees a genuine deferral window: batching low-severity fixes, or declining to pay continuous monitoring overhead for a minor issue, is the economically correct choice up to some finite T^* .

The caution this example makes visible is that the reasoning depends entirely on the assumption $\delta \approx 0$, an assumption that is easy to hold implicitly and rarely checked. A defect with active blast radius — an exploited vulnerability, a bug that corrupts data on every write — reintroduces the

compounding dynamics of §5.1, and the threshold that justified patience for an inert defect does not transfer to one that is actively growing. What makes software a useful illustration is precisely how easily this distinction is elided in practice: “we’ll fix it later” is a claim about δ , stated as though it were a claim about severity.

5.3 Infrastructure: Convexity and p

Bridges, water systems, rail networks, and electrical grids illustrate a regime dominated not by rapid compounding but by strongly convex repair cost. Replacing a handful of failing components is comparatively inexpensive; replacing a structure that has been allowed to deteriorate past a given point is not merely more expensive in proportion to the additional damage, but disproportionately so, since large-scale structural replacement carries costs — design, environmental review, extended service disruption, financing — that do not scale linearly with the physical extent of the repair.

This is a large- p regime, and the asymptotic estimate’s behavior under large p is exactly what the domain would predict: T_1 compresses toward 1 as p grows, meaning strongly convex repair cost punishes accumulated deviation even under modest growth rates. Infrastructure maintenance schedules that inspect and correct continuously, rather than waiting for visible failure, are an institutional recognition of this parameter specifically, largely independent of how fast any individual defect would otherwise compound.

5.4 Fiscal Reachability: Interaction of All Parameters

A fiscal system exhibits all four parameters simultaneously, and the framework is most informative here precisely because none of them can be treated as negligible in isolation. Auditing, reporting requirements, and administrative oversight constitute κ_M : standing costs paid whether or not any imbalance currently exists. Emergency intervention mechanisms — bailouts, restructuring, crisis-response facilities — carry a fixed initiation overhead κ_R , typically large relative to routine oversight, reflecting the political and institutional cost of invoking crisis powers at all. Structural imbalance, once established, tends to compound through interest-on-arrears dynamics and eroding institutional credibility, giving a genuine $\delta > 0$. And the cost of crisis-stage correction is strongly convex in the size of the imbalance being corrected, giving a large effective p .

Under this reading, T^* becomes a policy horizon: the point past which continuous, incremental stabilization is no longer merely prudent but strictly cheaper than deferred, discontinuous intervention, integrating over the accumulated effect of all four parameters at once. The interaction highlighted in §4.3 — that systems with a large nominal deferral tolerance (T_1 large, reflecting a currently well-functioning oversight regime with low κ_M relative to correction cost) are precisely the systems whose effective threshold erodes fastest under a rising δ — is the version of this framework’s warning most directly applicable to fiscal systems whose institutional credibility gives an appearance of tolerance that a rising rate of structural drift does not, in fact, support.

6 Relation to *Recursive Continuation*

This essay’s results refine rather than revise the apparatus of *Recursive Continuation*. The decomposition $R_t = M_t + \rho_t$ is compatible with every result of that essay’s Part II: the Complexity–Repair Theorem and Repair Dominance corollary are stated in terms of R_t as an undifferentiated capacity, and remain valid under any decomposition satisfying $R_t = M_t + \rho_t$, since nothing in their proofs depends on how R_t is internally constituted. What this essay adds is a finer question that the

coarser variable could not pose: not whether repair capacity keeps pace with maintenance burden in aggregate, but whether, conditional on it doing so, the burden is better paid continuously or discontinuously — a question invisible at the level of a single scalar R_t , and answered here by Proposition 3.4 and its asymptotic refinement.

The four failure modes of repair identified in that essay’s §10.4 — under-repair, over-repair, misrepair, and repair delay — receive a sharper treatment of the fourth under this essay’s apparatus. Repair delay was there described qualitatively, as an interval between detection and correction long enough for a fault to compound. This essay’s contribution is to make precise *when* such a delay is a failure and when it is not: delay is a failure mode only past T^* , and short of T^* the same delay is the economically correct exercise of the deferral window Proposition 3.4 establishes. Repair delay, in other words, is not intrinsically a defect of a recursive system; it becomes one only once a computable threshold, governed jointly by fixed overheads, compounding rate, and cost curvature, has been crossed.

Open Problems

1. **Non-constant δ and a .** This essay treats the compounding rate and per-step deviation as fixed. A fuller model would let δ_t itself depend on accumulated deviation, which is plausible in most of the domains treated in §6 (an infection’s growth rate is not generally constant; fiscal compounding accelerates with institutional credibility loss) and would likely remove the closed-form baseline T_1 in favor of a threshold defined only implicitly even at $\delta = 0$.
2. **Stochastic deviation.** Sections 2 through 5 treat a and δ as deterministic. A stochastic extension, in which deviation and compounding are random variables, would connect this essay’s threshold to the repair-entropy apparatus of *Recursive Continuation*, §10.3, since diagnostic uncertainty about the current value of d_t is plausibly a further cost this essay’s model does not yet price.
3. **Multiple correction channels.** This essay treats maintenance and repair as the only two available interventions. Real systems often have graduated intervention options between the two extremes (partial maintenance, minor versus major repair events), and the threshold structure under a continuum of intervention types, rather than a binary choice, is unresolved here.
4. **Empirical estimation of κ_M and κ_R .** As with $\bar{\rho}$ in *Recursive Continuation*, §14.5, this essay’s parameters are easier to define than to measure in a real system. Producing operational estimators for standing maintenance overhead and repair-initiation overhead in an actual institution remains open.

References

1. Richard E. Barlow and Frank Proschan. *Mathematical Theory of Reliability*. Wiley, 1965.
2. Andrew K.S. Jardine and Albert H.C. Tsang. *Maintenance, Replacement, and Reliability: Theory and Applications*, 2nd ed. CRC Press, 2013.
3. Ward Cunningham. “The WyCash Portfolio Management System.” *OOPSLA ’92 Experience Report*, 1992.

4. Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
5. Wassily Hoeffding. "Probability Inequalities for Sums of Bounded Random Variables." *Journal of the American Statistical Association*, 58(301), 1963.
6. Erik Hollnagel. *Safety-I and Safety-II: The Past and Future of Safety Management*. Ashgate, 2015.
7. Nancy G. Leveson. *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press, 2011.
8. Carmen M. Reinhart and Kenneth S. Rogoff. *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press, 2009.
9. American Society of Civil Engineers. *Failure to Act: Closing the Infrastructure Investment Gap for America's Economic Future*. ASCE, 2021.