

# Diagnostic Coordinates and the Time of Repair: A Correction to the Anti-Alignment Theorem

Flyxion

July 2026

## Abstract

History Before Function's Anti-Alignment Theorem shows that the most efficient compression of a history into an operator generically maximizes repair entropy, discarding the diagnostic surplus needed to repair the operator later. This note argues that the theorem, as stated, carries an unexamined assumption: that the trace evaluated at compression time and the trace available at repair time coincide. Four cases separate origin from availability. A folk recipe compressed from decades of trial and error never retains the coordinates at all. A theory-first legacy codebase retains them at design time and loses them passively through transmission. An error-correcting decoder retains coordinates sufficient for an assumed fault space and silently misapplies them, with full confidence, to faults outside that scope. A DRAM cell would lose its coordinates on a fixed physical schedule but does not, because a scheduled, active process continually regenerates them. Together these cases motivate diagnostic coordinate availability  $D(O, t; \mathcal{E})$ , evaluated at the time of repair and relative to a stated fault space, as the operative variable, with history-first compression, transmission decay, scope mismatch, and absence of maintenance standing as four distinct mechanisms by which  $D$  can fail rather than four different phenomena requiring separate theories. A further distinction, drawn from the difference between working within an accepted framework and adopting a new one, separates vertical repair, which recovers or extends coordinates already representable in an operator's diagnostic vocabulary, from horizontal repair, which requires a coordinate type that vocabulary never contained; the altitude recipe's eventual resolution is shown to be horizontal rather than a case of restored availability. The central claims are stated as conjectures, and a stripped-documentation test case designed to isolate origin from availability directly is proposed and left unrun.

## 1 The Theorem and Its Hidden Assumption

History Before Function proves an Anti-Alignment Theorem for its compression map  $C : H \rightarrow F$ , where a history  $H$  is compressed into an operator  $F$  governing future behavior. The theorem states, informally, that the most minimal retention of  $H$  consistent with  $F$ 's observed behavior weakly maximizes repair entropy, the log-count of fault locations consistent with a given symptom, and does so strictly whenever the discarded material would have excluded some fault hypothesis. Efficient compression, in other words, is generically the worst choice for later repairability, because execution and diagnosis are different objectives and compression is optimized for the former.

The theorem is stated against a single trace, evaluated once, at the moment of compression. That is an assumption, not a neutral simplification, and it is worth stating plainly: the theorem implicitly identifies the trace available for repair with the trace that was compressed in the first place. Call this the *synchrony assumption*. Under synchrony, an operator's repairability is fixed at birth by how it was compressed, and the natural corollary is a two-class picture: history-first operators, recipes, common-law traditions, biological adaptations, arise through unrecorded trial and error and are compressed-and-therefore-repair-poor from the start; theory-first operators, engineered systems with explicit derivations, retain their diagnostic coordinates because the derivation was never discarded to begin with.

The synchrony assumption is false in general, and two cases show why in complementary ways.

## 2 Two Cases That Separate Origin from Availability

### 2.1 The Altitude Recipe: History-First, Compressed at Birth

A recipe refined over generations at sea level absorbs observations such as dough rising too quickly, crust drying before the interior cooks, and proofing time varying with humidity into a handful of timings and ratios. The compression is optimal for execution: the recipe succeeds reliably within its original regime. It is pathological for diagnosis: moved to altitude, where boiling point and evaporation rate shift, the recipe fails, and nothing in the compressed operator distinguishes this failure from any other. The discarded distinctions were never retained: this is synchrony holding as stated, diagnostic poverty present from the moment of compression because the trace never contained the relevant coordinates.

### 2.2 The Legacy Codebase: Theory-First, Decayed by Transmission

A codebase built from an explicit design, with derivations, performance analyses, and documented rationale for every load-bearing constant, is theory-first in exactly the sense the two-class picture requires. Yet after years of refactoring, personnel turnover, and undocumented patches, the same codebase can present a repair situation indistinguishable from the recipe's: a magic constant nobody can currently explain, functionally

load-bearing, whose original justification existed but did not survive transmission. The operator itself, the running code, is unchanged. What changed is  $D(O, t)$ , the diagnostic coordinates available at the time repair is attempted, which decayed after the operator was built rather than being absent from it at birth.

This is the case the synchrony assumption cannot accommodate. Origin alone does not predict outcome: a theory-first operator can end up exactly as repair-poor as a history-first one, not because of how it was compressed, but because of what happened to its coordinates afterward.

### 3 Diagnostic Coordinates

**Definition 3.1** (Diagnostic Coordinate). *Given an operator  $O$  and a space of admissible fault hypotheses for  $O$ 's observed symptoms, a diagnostic coordinate is a retained distinction, in the sense already established by the Ecology of Distinctions' Axiom of Distinction, between two otherwise indistinguishable candidate histories that would have produced different present failures. Formally, if  $H_1$  and  $H_2$  are two histories compatible with  $O$ 's current behavior but would diverge under some future perturbation, a diagnostic coordinate is any retained information sufficient to separate  $H_1$  from  $H_2$  before that perturbation is observed.*

On this reading, a diagnostic coordinate is not a new primitive requiring its own theory. It is a distinction, in the corpus's existing sense, evaluated specifically over the space of candidate generating histories rather than over states in general. The recipe's altitude failure is precisely the case where no such distinction was ever retained between the histories consistent with sea-level success. The codebase's magic constant is the case where such a distinction existed and was subsequently lost.

**Definition 3.2** (Diagnostic Coordinate Availability).  *$D(O, t)$  is the number of diagnostic coordinates, in the sense above, retained and accessible for operator  $O$  at time  $t$ .*

### 4 First Correction: Temporal Availability

Repair entropy can be written explicitly as a function of which trace is consulted. Let  $T$  denote a retained trace,  $s$  an observed symptom,  $\mathcal{H}$  the space of candidate generating histories, and  $h \sim_T O$  the relation that  $h$  remains compatible with trace  $T$  and present operator  $O$ . Then

$$R(s | T) = \log |\{h \in \mathcal{H} : h \sim_T O, h \rightsquigarrow s\}|.$$

The Anti-Alignment Theorem, in this notation, is a claim about  $R(s | T_c)$ , repair entropy evaluated against the trace  $T_c$  retained at compression time. The synchrony assumption is the substitution  $T_r = T_c$ , identifying the trace available at repair time with the trace retained at compression. Under that substitution the theorem's conclusion is exact; once the substitution is dropped, the theorem gives  $R(s | T_c)$  while the quantity a repair process actually faces is  $R(s | T_r)$ .

**Proposition 4.1** (Synchrony Reduction). *If  $T_r = T_c$ , the corrected diagnostic-coordinate formulation reduces to the original Anti-Alignment Theorem.*

*Proof.* If  $T_r = T_c$ , every coordinate available at repair time is exactly one retained by the compression process, and no later degradation or regeneration changes the candidate-history partition. The repair entropy computed from  $T_r$  therefore equals the repair entropy computed from  $T_c$ , and the corrected formulation collapses to the compression-time theorem.  $\square$

This proposition is worth stating precisely because it fixes the corrected claim’s relationship to the original: not a rival result, but a strict generalization that recovers the original theorem exactly in the special case, repair immediately following compression, where the original theorem’s implicit assumption happens to hold.

**Proposition 4.2** (Compression Increases Repair Entropy). *Let  $T'$  be obtained from  $T$  by deleting a diagnostic coordinate that separates two candidate histories compatible with symptom  $s$ . Then  $R(s | T') \geq R(s | T)$ , with strict inequality whenever the deleted coordinate excluded at least one candidate history.*

*Proof.* Deleting a coordinate cannot remove any candidate history from the set compatible with the observed symptom; it can only merge previously distinguished histories. The set counted by  $R(s | T')$  therefore contains the set counted by  $R(s | T)$ , and taking logarithms preserves the inequality. If the deleted coordinate had excluded at least one candidate history, the containment is strict, so the entropy strictly increases.  $\square$

**Conjecture 4.1.** *Repair difficulty for an operator  $O$  at time  $t$  is governed by  $D(O, t)$ , diagnostic coordinate availability at the time repair is attempted, and not by whether  $O$ ’s origin was history-first or theory-first. Formally,*

$$\text{repair difficulty}(O, t) \propto \frac{1}{D(O, t)},$$

*with history-first compression standing as one mechanism, among others, by which  $D$  can be low, rather than as the defining condition for low repairability.*

Under this reading, the Anti-Alignment Theorem is not wrong; it is a special case of a more general claim, correct precisely when  $D(O, t_{\text{compression}}) = D(O, t_{\text{repair}})$ , which holds automatically for operators repaired shortly after compression but is not guaranteed by the compression process itself. The theorem’s real content, on this reading, was never about history-first origin per se. It was an early, correct observation about one specific mechanism, execution-optimized compression, that reliably produces low  $D$  at the moment of compression, stated in a form that did not yet distinguish it from availability at arbitrary later times.

A passive-decay model makes the codebase case concrete. Let  $D(O, t) = D(O, t_0)e^{-\lambda(t-t_0)}$ , where  $\lambda$  aggregates personnel turnover, documentation loss, refactoring drift, and institutional forgetting.

**Proposition 4.3** (No Fixed Advantage Survives Unbounded Decay). *For any history-first operator  $O_h$  with low initial diagnostic availability  $D(O_h, t_0)$ , and any theory-first operator  $O_f$  with high initial diagnostic availability  $D(O_f, t_0)$  and decay rate  $\lambda > 0$ , there exists a finite later time  $t^*$  such that  $D(O_f, t^*) \leq D(O_h, t_0)$ .*

*Proof.* Since  $\lambda > 0$ ,  $D(O_f, t_0)e^{-\lambda(t-t_0)} \rightarrow 0$  as  $t \rightarrow \infty$ , so for any fixed positive threshold  $D(O_h, t_0)$  there exists a finite  $t^*$  at which  $D(O_f, t^*) \leq D(O_h, t_0)$ .  $\square$

This proposition is deliberately modest, and it is worth being explicit about the scope of what it shows: it is a generic fact about exponential decay toward zero, true of any positive quantity decaying at any rate  $\lambda > 0$ , and it does not by itself distinguish anything specific to theory-first operators beyond the presence of decay. What it establishes, precisely because it needs so little, is that no fixed initial advantage in  $D$  is protective against unbounded neglect; origin sets  $D(O, t_0)$ , but nothing about origin bounds  $\lambda$ .

## 5 A Test Case

The clearest way to separate origin from availability empirically is a case where they are pulled apart by construction rather than by accident. Consider a theory-first operator, an engineered component with a complete design derivation, whose documentation is then deliberately removed after the fact: comments stripped for intellectual-property reasons, the derivation withheld from the shipped artifact, the operator itself left completely unchanged. This differs from the codebase case in an important way: transmission decay there is gradual and unintentional, confounded with genuine forgetting across personnel and time. Deliberate stripping isolates the variable directly, theory-first origin held fixed,  $D(O, t)$  manipulated directly and independently of any historical compression process.

**Open Question 5.1.** *If repairability of such a component drops to match the recipe's failure profile once its documentation is stripped, that is strong evidence that  $D(O, t)$  is the operative variable and origin is not independently protective. If repairability instead remains higher than the recipe case despite the stripping, something about theory-first origin is doing work the diagnostic-coordinate reframing has not yet captured, and the correction proposed here is incomplete.*

This question is left open, and it names one failure mode, silent absence, where repair simply cannot proceed because  $D(O, t)$  has dropped too low to distinguish any fault hypothesis at all. There is a second, sharper failure mode that the recipe and codebase cases do not by themselves reveal, and it deserves separate treatment because it is not a matter of degree along the same axis.

## 6 Second Correction: Fault-Space Dependence

Two distinct things can go wrong when a coordinate set meets a dynamics it is asked to support, and the paper so far has vocabulary for only one of them.

**Definition 6.1** (Operational Sufficiency). *A coordinate set  $T$  is operationally sufficient for a dynamics  $F$  if  $F$  can be evaluated using only information contained in  $T$ .  $T$  is operationally insufficient for  $F$  when evaluating  $F$  requires distinctions absent from  $T$ .*

Wilson’s discussion of levels supplies a case of the first kind, refusal rather than silent error. A macroscopic description, a chair, characterized without its constituent particles’ spin and position, cannot be evaluated by the equations of quantum mechanics at all: those laws require degrees of freedom the description does not carry, and no output is produced, correct or otherwise, until the missing information is supplied [4]. This is operational insufficiency in the strict sense of the definition above:  $F$  does not run.

The Hamming decoder is a different failure entirely. Its coordinates are operationally sufficient for computing a syndrome from any received word, single-bit fault or not; the syndrome function evaluates without complaint on inputs its designers never anticipated. What is insufficient is not the ability to run  $F$ , but the ability of  $F$ ’s output to discriminate the actual fault from a merely consistent one. The decoder does not refuse; it answers confidently and wrongly. Distinguishing these two modes, a dynamics that cannot be evaluated at all against one that evaluates but misidentifies its input, matters because they call for different remedies: the first is fixed by supplying the missing information before evaluation is attempted, the second only by adding coordinates capable of separating cases that current ones conflate, which is exactly the scope-failure problem this section develops next.

**Definition 6.2** (Scoped Diagnostic Sufficiency). *An operator’s diagnostic coordinates at time  $t$  are scoped-sufficient for a fault space  $\mathcal{E}$  if they discriminate every pair of fault hypotheses within  $\mathcal{E}$ , and merely scoped if  $\mathcal{E}$  is a strict subset of the space of faults the operator will actually encounter.*

The Hamming case shows that  $D$  is not adequately described by a single count of retained coordinates. Let  $\mathcal{E}_{\text{design}}$  be the fault space a repair operator’s coordinates were built to resolve, and  $\mathcal{E}_{\text{actual}}$  the fault space actually encountered. Diagnostic sufficiency requires an injective syndrome map  $\sigma : \mathcal{E} \rightarrow \Sigma$ , where  $\Sigma$  is the set of observable diagnostic syndromes; the availability quantity is more properly written  $D(O, t; \mathcal{E})$ , indexed by fault space rather than treated as a bare count. It is useful to make this indexing precise as a subset relation rather than only as a scalar.

**Definition 6.3** (Diagnostic Profile). *Let  $P(\mathcal{E})$  denote the set of fault distinctions available in fault space  $\mathcal{E}$ . The diagnostic profile of operator  $O$  at time  $t$  is  $\Pi_D(O, t) \subseteq P(\mathcal{E})$ . The operator is diagnostically complete for  $\mathcal{E}$  if  $\Pi_D(O, t) = P(\mathcal{E})$ , and diagnostically partial relative to  $\mathcal{E}$  if  $\Pi_D(O, t) \subsetneq P(\mathcal{E})$ .*

The Hamming decoder’s profile is then  $\Pi_D(O, t) = P(\mathcal{E}_{\text{single}}) \subsetneq P(\mathcal{E}_{\text{actual}})$ , complete for single-bit faults and partial for the actual fault space once double-bit faults are possible, which states the scope-failure result of this section as a subset relation rather than only as an informal comparison. The term *diagnostically emergent* suggested itself here by analogy to Wilson’s weak emergence, but is deliberately not adopted: Wilson’s proper-subset account earns autonomy through a specific further argument, that the higher-level entity’s distinct powers make it non-identical to its base by Leibniz’s Law, and nothing in

a decoder's coordinate profile does comparable work. *Diagnostically partial* claims only what the definition above establishes.

**Proposition 6.1** (Scope Failure). *If there exist  $e_1 \in \mathcal{E}_{\text{design}}$  and  $e_2 \in \mathcal{E}_{\text{actual}} \setminus \mathcal{E}_{\text{design}}$  such that  $\sigma(e_1) = \sigma(e_2)$ , then no deterministic repair procedure using only  $\sigma$  can distinguish  $e_1$  from  $e_2$ .*

*Proof.* A deterministic repair procedure using only  $\sigma$  receives the same input for  $e_1$  and  $e_2$  and must therefore return the same output for both. If the correct repairs for  $e_1$  and  $e_2$  differ, the procedure misrepairs at least one of them.  $\square$

**Remark 6.1.** *The deterministic case is what the circuit example needs, since a syndrome decoder is a deterministic function of the syndrome. A randomized procedure with access to a prior over fault likelihoods could in principle do better than chance even without deterministic distinguishability; the proposition above does not rule this out; the claim proved here is only that no repair choice can be certified by  $\sigma$  alone, not that no procedure whatsoever can outperform pure guessing.*

Error-correcting memory gives a clean, checkable instance of this scope failure. A Hamming(7,4) code retains exactly three syndrome bits, sufficient to distinguish any single-bit-error hypothesis,  $\mathcal{E}_{\text{design}}$ , from the transmitted codeword.  $D(O, t; \mathcal{E}_{\text{design}})$  is high: the coordinates needed to diagnose the design-assumed fault space are present, retained, and mechanically checkable on every read [3].

For a single-bit fault  $e_i$ , the syndrome points to position  $i$  and correction restores the transmitted codeword. For a double-bit fault, the syndrome computed from the same coordinates is generally nonzero and coincides with the syndrome of some single-bit fault  $e_k$ ; the decoder corrects position  $k$  with full confidence, and the resulting word generally differs from the original by a three-position deviation rather than a correction. Because the result satisfies every parity check the retained coordinates are capable of running, it is admissible by every test available to the repair operator, and it is not the object that was actually transmitted.

**Proposition 6.2** (Confident Error). *If an out-of-scope fault shares a syndrome with an in-scope fault, and the repair operator always corrects in-scope syndromes, the operator will report success on at least some out-of-scope faults.*

*Proof.* Let  $e_o$  be an out-of-scope fault and  $e_i$  an in-scope fault with  $\sigma(e_o) = \sigma(e_i)$ . By assumption the operator interprets this syndrome as evidence for  $e_i$  and applies the corresponding correction. Since every check available to the operator is a function of the same syndrome, and the corrected object is evaluated only against those checks, the operator has no internal basis for refusing the repair, and reports success even though the actual fault was not the one repaired.  $\square$

This sharpens the first correction rather than replacing it: an operator can have  $D(O, t; \mathcal{E}_{\text{design}})$  high while having no diagnostic capacity whatsoever relative to  $\mathcal{E}_{\text{actual}} \setminus \mathcal{E}_{\text{design}}$ , with no signal internal to the operator marking the boundary between the two. The practical circuit response is to add redundancy beyond the original design distance, an overall parity bit converting single-error-correction into single-error-correction-with-double-error-detection, which still cannot repair the larger fault but can at least refuse to

certify a wrong repair as successful [3]. That is a third possible state for an operator to be in, beyond simply high or low  $D$ : diagnostically silent about its own scope, confidently wrong, or honestly reporting the limits of what it can currently distinguish. Which of these three an operator occupies is not visible from repairability alone, since silent miscorrection and successful repair look identical from outside until the wrong repair's consequences surface later.

The scoped-sufficiency result also corrects an assumption that has been running informally underneath the codebase case since it was first introduced: that more retained coordinates is simply better. That was never quite the claim, and it is worth stating the correct version explicitly.

**Proposition 6.3** (Diagnostic Autonomy). *An operator need not retain every distinction present in its generating history to remain repairable. It need only retain those distinctions that discriminate faults within the relevant fault space  $\mathcal{E}_{actual}$ ; additional retained distinctions may increase descriptive fidelity to that history without increasing repairability.*

This is stated as a proposition rather than proved from anything already established, and its justification is a motivated analogy rather than a derivation: Wilson's weak-emergence account of ordinary objects treats a higher-level entity's autonomy as arising from retaining a proper subset of its physical base's causal powers, not the full set, with additional lower-level powers contributing to descriptive completeness without contributing to the higher-level entity's own causal profile [4]. The parallel suggested here is that repairability, like Wilson's causal autonomy, tracks selective sufficiency rather than maximal retention, but the two claims rest on different grounds, hers on Leibniz's Law applied to distinct power profiles, this one on the definition of scoped diagnostic sufficiency above, and the analogy should be read as motivating the proposition rather than establishing it.

## 7 Third Correction: Horizontal Repair

Everything so far treats repair as movement within an already-fixed representational scheme: coordinates are present or absent, sufficient or scoped, but in every case they are drawn from a vocabulary, ratios and timings for the recipe, syndrome bits for the decoder, that was fixed at the moment the operator was compressed. A different part of Wilson's talk, its opening discussion of progress in metaphysics and physics rather than the weak-emergence material used above, distinguishes vertical progress in a discipline, working out the internal consistency and consequences of accepted framework assumptions, from horizontal progress, jumping to an entirely different set of framework assumptions [4]. That distinction is Wilson's own, stated about theoretical inquiry, not about repair; what follows transposes it, treating it as a motivated parallel for a first-order distinction about operators rather than as an application of anything Wilson herself argues about engineering. Repair admits the analogous split, and the two cases already on the table divide along it in a way worth making explicit.

**Definition 7.1** (Vertical and Horizontal Repair). *A repair is vertical when diagnosis proceeds by recovering or extending coordinates within an operator’s existing representational scheme: the same type of syndrome, the same kind of distinction, applied more completely or restored after decay. A repair is horizontal when successful diagnosis requires a coordinate of a kind not previously represented in the operator’s diagnostic vocabulary at all, so that the fault space itself must be redescribed rather than merely covered more fully by existing means.*

The SECDED extension of the earlier Hamming case is vertical by this definition. Adding an overall parity bit to convert single-error-correction into single-error-correction-with-double-error-detection stays entirely within the existing representational scheme: one more syndrome bit of the same kind as the three already present, extending  $\mathcal{E}_{\text{design}}$  without introducing any new type of coordinate. The scheme gets more of what it already had.

The altitude recipe’s eventual resolution is horizontal, and it is worth being precise about why recovery is not the right description of what happens. Nothing about the original compression ever separated sea-level histories from altitude histories; there was no coordinate to lose, decay, or fail to recover, in contrast to the codebase case, where the original derivation existed and could in principle be reconstructed from old records or a former engineer’s memory. When a baker eventually diagnoses the altitude failure, what gets introduced, boiling point, atmospheric pressure, evaporation rate, is not a restored fragment of the original compression. It is a coordinate of a kind the recipe’s representational vocabulary, ratios and timings, never contained and could not have expressed. The repair does not recover  $D(O, t; \mathcal{E})$  toward some earlier value; it enlarges the space of representable coordinates itself.

**Proposition 7.1** (Horizontal Repair Is Not Availability Restoration). *Let  $C(O, t)$  denote the set of coordinate types representable in operator  $O$ ’s diagnostic vocabulary at time  $t$ . Vertical repair operates entirely within a fixed  $C(O, t)$ , increasing  $D(O, t; \mathcal{E})$  for  $\mathcal{E}$  expressible using coordinates already in  $C(O, t)$ . Horizontal repair requires  $C(O, t') \supsetneq C(O, t)$  for some later  $t'$ : a coordinate type not representable in  $C(O, t)$  becomes representable, and only then can the corresponding fault space be diagnosed at all.*

*Proof.* If a fault hypothesis distinguishing  $e_1$  from  $e_2$  requires a coordinate type absent from  $C(O, t)$ , no syndrome map  $\sigma$  constructible from  $C(O, t)$  can separate  $e_1$  from  $e_2$ , by the same argument as the Scope Failure proposition above, since  $\sigma$ ’s domain is fixed by the representable coordinate types available to construct it. Distinguishing  $e_1$  from  $e_2$  therefore requires a coordinate type outside  $C(O, t)$ , which is exactly  $C(O, t') \supsetneq C(O, t)$  for the  $t'$  at which that type becomes available.  $\square$

This proposition is the least settled claim in the paper, and it is worth flagging why: everything else here sharpens or corrects a single existing theorem within one representational scheme, and the diagnostic-coordinate vocabulary was deliberately built to avoid new ontology, grounding itself in the Ecology of Distinctions’ existing apparatus rather than introducing a fresh primitive.  $C(O, t)$ , the set of representable coordinate types, is not obviously reducible to that apparatus in the same way; whether horizontal repair can be given the same treatment as the rest of this note, as a special case of an existing

recoverability or distinction-theoretic result, or whether it genuinely requires new formal machinery this corpus does not yet have, is left open rather than resolved here.

## 8 Active Regeneration as a Fourth Case

The codebase case showed  $D(O, t)$  decaying passively after compression, coordinates present at design time and lost through ordinary transmission. Dynamic memory circuits show the opposite discipline is possible: a DRAM cell's stored charge decays according to  $V(t) = V_0 e^{-(t-t_{\text{last}})/\tau}$ , and the cell remains admissible only while  $t - t_{\text{last}}$  stays below a threshold set by  $\tau$  and the minimum distinguishable voltage [3]. Left alone, such a cell would follow the codebase's trajectory exactly, diagnostic and even representational capacity draining away on a fixed physical timescale. What prevents this is not a more efficient initial compression but a scheduled, active process, refresh, that rewrites the cell before decay crosses the threshold, resetting  $t_{\text{last}}$  and regenerating the condition for the cell's own future readability.

This gives the corrected claim a fourth case to sit alongside the recipe, the codebase, and the miscorrecting decoder: an operator whose  $D(O, t; \mathcal{E})$  would decay like the codebase's but is instead held constant by continuous, deliberate reinvestment, rather than by a one-time compression choice at any point in its history.

**Proposition 8.1** (Refresh as Coordinate Regeneration). *If a memory cell is refreshed at intervals  $\Delta t < \tau \log(V_0/V_{\min})$ , its readability condition is maintained indefinitely under the exponential-decay model.*

*Proof.* After each refresh the voltage is restored to  $V_0$ . The minimum voltage before the next refresh is  $V_0 e^{-\Delta t/\tau}$ , and  $\Delta t < \tau \log(V_0/V_{\min})$  gives  $V_0 e^{-\Delta t/\tau} > V_{\min}$ , so the cell never falls below threshold. The same argument applies after every subsequent refresh.  $\square$

Repairability, on this reading, is not fixed by an operator's origin, nor is it simply a quantity that decays; for at least one real physical system, it is actively maintained against a known decay process by treating diagnostic availability as something requiring ongoing work rather than a property to be gotten right once.

## 9 Conclusion

The correction proposed in this note now rests on four cases, distinguished formally by writing diagnostic availability as  $D(O, t; \mathcal{E})$  rather than as a bare count. The recipe has low  $D(O, t_c; \mathcal{E}_{\text{actual}})$  because the relevant distinctions were never retained. The codebase has high  $D(O, t_c; \mathcal{E}_{\text{actual}})$  at compression but declining  $D(O, t_r; \mathcal{E}_{\text{actual}})$  at repair, through passive transmission decay. The Hamming decoder has high  $D(O, t; \mathcal{E}_{\text{design}})$  but effectively zero  $D(O, t; \mathcal{E}_{\text{actual}} \setminus \mathcal{E}_{\text{design}})$ , a scope failure rather than an absence. The DRAM cell has a naturally decaying coordinate condition prevented from mattering by active regeneration.

What began as a single temporal correction, that repair time and compression time need not coincide, is now four distinct claims rather than one: that the two times must be separated; that diagnostic coordinates are only meaningful relative to a stated fault space, not as a bare count; that availability can be produced, lost, or restored over time rather than simply decaying monotonically from a value fixed at origin; and that some repairs cannot be described as availability restoration at all, because they require a coordinate type the operator’s representational scheme never contained. Origin explains none of the four cases’ outcomes on its own. The altitude recipe, introduced early as an instance of coordinates simply never retained, turns out on closer inspection to motivate the fourth claim rather than only the first: its failure is a matter of absent  $D$ , but its eventual repair, if it comes, is horizontal rather than a recovery of anything that was ever there to recover. A further question these claims raise, and one this note does not attempt to answer, is how  $D(O, t; \mathcal{E})$  depends jointly on the three quantities this note has now distinguished: representable vocabulary  $C(O, t)$ , temporal availability, and fault space. The DRAM and Hamming cases both hold vocabulary fixed while varying time or scope respectively; the altitude recipe is the only case among the four where vocabulary itself is what changes, and that asymmetry is worth taking seriously rather than treating as incidental. It is tempting to write  $D(O, t; \mathcal{E}) = A(O, t) \cdot S(O, \mathcal{E})$  and ask when such a factorization holds, but committing to a multiplicative form this early would assume more structure than the four cases actually establish. The more honest open question is simply

$$D(O, t; \mathcal{E}) = f(C(O, t), A(O, t), \mathcal{E})$$

for some dependency  $f$  not yet characterized, together with the narrower question of whether  $C(O, t)$  functions as a bound on the other two arguments, as the discussion of scope in Section 6 suggested, or as a fully independent third factor. Both remain open, and settling either is left for a paper this note has not yet become. The stripped-documentation test proposed above remains the cleanest way to isolate origin from availability directly, and it is still unrun.

## A Diagnostic Coordinates as Partitions of Candidate-History Space

The repair-entropy argument in Section 3 is already, implicitly, an argument about partition coarsening; this appendix states it as one explicitly, since doing so shows precisely why scope failure and coordinate loss are instances of a single underlying phenomenon rather than two unrelated failure modes.

**Definition A.1** (Induced Partition). *For a retained trace  $T$ , let  $\mathcal{P}(T)$  be the partition of the candidate-history space  $\mathcal{H}$  in which  $h, h' \in \mathcal{H}$  lie in the same block if and only if  $T$  cannot distinguish them,  $h \sim_T O \Leftrightarrow h' \sim_T O$  for the operator  $O$  in question.*

**Definition A.2** (Coarsening).  $\mathcal{P}(T') \preceq \mathcal{P}(T)$  when  $\mathcal{P}(T')$  is coarser than  $\mathcal{P}(T)$ , that is, every block of  $\mathcal{P}(T)$  is contained in some block of  $\mathcal{P}(T')$ . If  $T'$  is obtained from  $T$  by deleting

a diagnostic coordinate,  $\mathcal{P}(T') \preceq \mathcal{P}(T)$ , since deleting a coordinate can only merge previously separated blocks.

**Proposition A.1.** *Compression Increases Repair Entropy (Section 3) is a corollary of coarsening:  $R(s | T)$  is the logarithm of the size of the block of  $\mathcal{P}(T)$  containing the histories compatible with symptom  $s$ , and coarsening a partition cannot decrease, and generically increases, the size of the block containing any fixed element.*

The scope-failure result of Section 6 admits the same reading at one remove: a fault space  $\mathcal{E}$  supplies a second partition, this one on fault hypotheses rather than histories, and scoped sufficiency is exactly the requirement that the partition induced by an operator's coordinates refine the partition an adequate diagnosis of  $\mathcal{E}$  would require.

**Definition A.3** (Diagnostic Profile Lattice). *Order diagnostic profiles by inclusion,  $\Pi_1 \leq \Pi_2 \iff \Pi_1 \subseteq \Pi_2$ . The Hamming(7,4) decoder occupies the node  $\Pi_D(O, t) = P(\mathcal{E}_{\text{single}})$ ; the SECDED extension occupies a strictly larger node, adding double-error detection without full double-error correction; complete diagnosis relative to  $\mathcal{E}_{\text{actual}}$  is the top element of the lattice restricted to profiles achievable without enlarging  $C(O, t)$ .*

This lattice gives Section 6's informal comparison, "high relative to  $\mathcal{E}_{\text{design}}$ , absent relative to  $\mathcal{E}_{\text{actual}} \setminus \mathcal{E}_{\text{design}}$ ," a precise ordering rather than only a verbal one. It is deliberately not extended further here. A natural next step would be to posit an operator  $\Phi$  mapping  $C(O, t)$  to the diagnostic profiles it makes achievable,  $\Pi_D(O, t) = \Phi(C(O, t))$ , but  $\Phi$  introduced this way would carry no stated properties, no domain or codomain structure beyond the bare claim that it exists, and defining it would give the appearance of formalizing the relationship between vocabulary and profile without actually constraining it. The honest statement of that relationship is the open question closing the main text, that  $\Pi_D(O, t)$  is bounded by  $C(O, t)$  in some manner not yet characterized, and this appendix stops at the point where further notation would outrun what has actually been established.

## References

- [1] Flyxion. *History Before Function: Compression as the Origin of Operators*. Admissibility Program, 2026. <https://standardgalactic.github.io/alphabet/example/history-before-function.pdf>
- [2] Flyxion. *The Ecology of Distinctions*. Admissibility Program. [https://standardgalactic.github.io/spherepop/textbook/The\\_Ecology\\_of\\_Distinctions.pdf](https://standardgalactic.github.io/spherepop/textbook/The_Ecology_of_Distinctions.pdf)
- [3] Flyxion. *The Continuation Geometry of Circuits*. Independent Researcher, 2026. [https://standardgalactic.github.io/robotics/continuation\\_geometry\\_of\\_circuits.pdf](https://standardgalactic.github.io/robotics/continuation_geometry_of_circuits.pdf)
- [4] Wilson, J. *The Emergence of Ordinary Objects*. Talk presented at Metaphysics Within and Without Physics, Western University, London, ON, June 7–8, 2014. Supported by the Social Sciences and Humanities Research Council (SSHRC), the Rotman Institute of Philosophy, the Department of Philosophy, the Faculty of Arts and Humanities, and Research Western. <https://youtu.be/SNnb7wW0woQ>