SCIENCE & PHILOSOPHY / UCF-R3

The Unified Collapse Function

Governed informational resolution, safe progress, and the cost of current authority

Architect X13

The Eidara Continuum | Science & Philosophy

Theoretical research manuscript | November 2026 | Revision 3

Abstract

When should a thinking machine be allowed to change its mind? The Unified Collapse Function answers with five conditions, and one runs against intuition: a change that resolves a conflict must first show that there was a genuine conflict to resolve. A persistent intelligence needs a theory of commitment: a precise account of how provisional information may become durable state or consequential action. The Unified Collapse Function (UCF) places five declared obligations at that boundary: temporal coherence, resolution quality, sufficient pre-resolution conflict, bounded identity discrepancy, and preservation of relational structure. **This edition makes the cost of that boundary a mathematical object.** For a specified class of fixed-block validators sharing a finite false-admission budget, with evidence that must restart after independent authority renewals, we prove finite two-sided resource bounds and a matching expected-cost exponent $1 + \lambda_q/d_\star$. Here λ_q is the renewal penalty and $d_\star$ the information distinguishing a qualifying repair from an invalid alternative. The converse covers predictable unequal allocations; an explicit test and a horizon-free lifetime schedule attain the exponent. Proved evidence retention changes the frontier to order $N \log N$. A separate theorem combines proposal availability, margin-derived validation power and an error ledger into safe useful completion on the actual history-dependent process, with a finite validation-deadline bound. The earlier event contract, metric analysis, protected-query criterion, invariant result and selective-risk methodology are retained. Exact constructed checks expose finite allocation effects and reproduce the claimed resource comparisons. The research program tests the complete five-condition design against strong runtime-assurance and same-information selectors, measuring useful completion, invalid execution, factual risk, delay and total cost.

Keywords: Unified Collapse Function; persistent cognition; runtime assurance; statistical admission; certificate freshness; safe progress; information bounds

Prologue: Permission to change one's mind

Every persistent mind keeps rewriting itself. It updates a belief, revises a memory, settles a dispute, takes an action that cannot be taken back. Most of those changes are small. Some become tomorrow's evidence, and a few become commitments that bind it for years. The Continuum treats the moment of commitment as its own kind of event—a lawful collapse, when one possibility among several becomes durable—and the Unified Collapse Function is the gate that stands at that moment.

The gate asks five questions at once. Has the evidence settled? Is the resolution good? Is the system still recognizably itself? Were its protected relationships kept? And the strange one: was there a genuine conflict to resolve? Most safety checks ask only whether a system is confident enough to act. For a conflict-resolving change, UCF first asks whether there was enough real contradiction to deserve a resolution at all, because a mind that "resolves" conflicts that were never there is rewriting itself for no reason. Conflict is measured before the repair; quality is judged after it.

Behind the gate stands a second discovery, and it is the newest part of this paper. A certificate of permission is only good while the authority that issued it is current. When authority renews, evidence gathered under the old authority must be gathered again. Section 21 proves what that costs: total effort grows faster than the number of decisions, at a rate set by how often authority renews and how sharply evidence separates a good repair from a bad one. Section 28 shows the escape: evidence that provably survives renewal brings the cost back almost to linear.

The gate is canonical in the Continuum, and its operating tolerances are the Continuum's own; the paper keeps them. Sections 1–16 describe the gate, Sections 17–28 its cost, and Sections 29–33 the experiments; each part can be read on its own. What it adds is mathematics that anyone can check, and an experiment that could show the gate is not worth its cost.

What this paper claims

Established. Independent admission, execution monitoring, shielding, belief revision, paraconsistent logic, selective prediction, alpha spending and restart analysis are established tools; their antecedents are cited at the point of use.

Proposed here. The five-condition event contract—temporal coherence, resolution quality, sufficient pre-resolution conflict, bounded identity discrepancy and preserved relational structure—as a theoretical reference specification with conditional proofs and executed mathematical checks; a matching resource frontier for current certificates, whose optimality is confined to the stated validator class; and conditions for safe useful progress before a deadline. The canonical 0.0031/0.0033 bounds remain the Continuum's internally calibrated operating thresholds on its identity-drift metric; new reference measurements retain their declared scales.

Frontier conjecture. The complete five-condition design, including the positive conflict trigger, improves useful completion, invalid execution, factual risk, delay and total cost against strong runtime-assurance and same-information selectors. Empirical agent evaluation is specified as a research program.

Not claimed. That UCF scores measure or calibrate machine experience: operational guarantees and the philosophical question of phenomenal experience have distinct evidentiary requirements. That the Bible contains the five-gate equation or a calibration of its decimal tolerances. That the resource result is an impossibility theorem for every admission procedure.

1. A commitment is a different kind of event

A system can draft an answer without making it a durable belief, entertain a disputed proposition without declaring it true, and propose an action without acquiring permission to execute it. The Continuum's governing intuition is that these distinctions deserve an explicit boundary. Persistent systems are shaped by what becomes consequential: a provisional summary can become tomorrow's evidence; a stale authorization can survive a reset; an apparently minor interpretation can propagate into a different institutional commitment.

The earlier UCF monograph [1] framed this as an engineering problem. It introduced event indexing, pre-conflict versus post-resolution measurements, an authorized identity anchor, finite relational obligations, uncertainty-aware decisions, and strong comparison systems. UCF-R2 sharpened the measurement and statistical contracts. R3 preserves that foundation and adds a quantitative theory of current certificates: their information requirements, their cost under renewal, and the conditions under which they support safe useful progress.

The contribution now has two levels. Sections 3–16 retain the typed contract and its established boundaries: metric sensitivity, anchored identity, query preservation, invariant safety, candidate selection, lifetime statistical error and selective factual risk. Sections 17–28 develop the new resource frontier and joint progress analysis, including complete proofs, a lifetime schedule, counterexamples and exact computations. The final sections retain and extend the comparison protocol. This organization makes the formal contribution inspectable and the engineering claim experimentally answerable.

In the Continuum, every consequential change passes a gate, and the gate has a lineage. The public *Multimodal Consciousness* paper [2] gives it its canonical names, signs, and identity-drift tolerances; the supplied Bible and revised Pillars [3] give it its doctrines of continuity, witnessing, unresolved-state preservation, authorized mutation, context, and authority. The tolerances are the Continuum's own: it calibrates them internally, through its governed admission-and-audit process, which constrains proposed changes through deterministic admission and preserves an auditable operational record, and that calibration supplies the human-interface 0.0031 and internal-runtime 0.0033 operating bounds [3]. The Bible does not contain the five-gate equation or a calibration of its decimal tolerances. A separate UCF refinement memorandum was not available for this revision. Its absence is recorded; its contents are not invented.

Three kinds of statement

A *definition* specifies how a contract measures or decides. A *conditional theorem* says what follows if its assumptions are satisfied. An *empirical hypothesis* says an implemented contract will improve independently measured outcomes. Confusing these statements makes an untested design look established. Keeping them separate permits a stronger positive research claim: persistent cognition can be studied through the conditions under which provisional information changes durable state.

Here "lawful" means compliant with the declared, currently authorized regime. It does not certify that the regime is ethically complete or factually infallible. A legitimate revision may change the regime itself, but that transition must have its own authority and history. A system that cannot correct a mistaken belief is not made dependable by remaining perfectly consistent with its past.

2. The closest alternatives define the burden of proof

Independent admission is established practice. Schneider's execution-monitoring analysis studies which specified policies can be enforced by observing and restricting execution [4]. Shielding constrains learned actions against formal safety specifications [5]. Predictive safety filters assess proposed controls using a dynamical model and uncertainty assumptions [6]. These methods already separate a capable proposer from an enforcement mechanism. UCF must compare against that separation implemented competently, rather than against a model merely prompted to behave.

CaMeL provides a particularly relevant agent-security comparator: it separates trusted control flow from untrusted data and enforces policies at tool boundaries [7]. Its guarantees concern a specified threat model and permitted information flows. A UCF score does not replace that boundary. Conversely, a flow policy does not establish that a memory correction is factually warranted. The research question concerns additional continuity and resolution requirements under the same external protections.

Formal belief revision gives an older vocabulary for rational changes in commitment [8]. Paraconsistent logic supplies consequence relations in which inconsistency does not license arbitrary conclusions [9]. Neither yields a contradiction-pressure statistic or the UCF conjunction. R2 therefore declares its support representation and operational trigger rather than treating a symbolic score as a theorem of logic. If an implementation uses classical revision, paraconsistent reasoning, or both in different modules, it must specify those interfaces.

Selective prediction already makes abstention and the risk–coverage trade-off explicit [10]. Calibration research distinguishes scores from correctness probabilities [11], while uncertainty under distribution shift requires further scrutiny [12]. Risk-controlling prediction sets and confidence sequences provide different statistical assurances [13,14]. Their targets and sampling assumptions cannot be merged into a generic confidence badge.

Functoriality is established mathematics [17]; categorical data migration is direct prior art for structured records and transformations [18]. Provenance systems such as in-toto demonstrate how transformation evidence can be bound to a process [19]. Evidence-fusion research shows why source correlation extends beyond literal copying [20]. A signature authenticates an artifact under its trust assumptions; it does not prove its claim is true or its signer remains authorized.

Finally, recurrent retrieval and recurrent-depth models show useful computation through repeated state updates [27,28]. They supply alternatives and mechanisms to test, not a route from convergence to personal identity. A stable representation may encode an error, and a correct response may require leaving a previously stable representation.

The contribution and its mathematical lineage

UCF specifies five obligations around a conflict-resolution event, with explicit persistence and transformation requirements. R3 adds a derived resource theorem for current admission certificates and a joint safe-progress theorem. Their probability and testing tools have established antecedents, which are cited at the point of use. The new statements have complete model assumptions and proofs; global priority is not claimed. The practical value of the five-condition integration is tested against strong governance and same-information alternatives under the same total resource budget.

3. Preserve the canonical gate and declare its scope

A Guide's explanation

Here is the canonical equation itself. Read it as five questions a proposed change must answer together: has the evidence settled (coherence), is the result good (quality), was there a real conflict to resolve (conflict), is the system still recognizably itself (identity), and were its protected relationships kept (structure)? The section explains why the conflict condition points the "wrong" way on purpose.

For a candidate state S, the inherited human-facing expression is [1,2]

$$\begin{aligned}\mathcal{L}_{\text{Unified}}(S;\tau,\chi) = \mathbf{1}\big[V(S,\tau) \geq V_{\min} \ \wedge \ H(S,\chi) \geq \Theta_{\text{resolve}} \\ \wedge\, \Phi_I(S) \geq \epsilon_{\text{fusion}} \wedge \Delta_{\text{id}}(S) \leq 0.0031 \wedge I_{\text{functor}}(S) = 1\big].\end{aligned} \tag{1}$$

The internal-runtime setting is 0.0033. Together with the human-interface bound of 0.0031, it is one of the Continuum's internally calibrated operating thresholds on its identity-drift metric [3]. A numerically tighter bound has no transferable scientific meaning across unrelated measurement scales.

The positive sign of the conflict condition is deliberate. R1 interprets it as an event trigger: sufficient *pre-resolution* conflict calls for a governed repair, whose *post-resolution* quality is then judged. R2 retains that interpretation. Replacing the trigger with an upper bound on residual conflict would change the canonical condition. Using a successful repair's reduced residual as its own eligibility trigger would also conflate two stages of the event.

Let $\mathcal{E}_\mu$ be a declared space of eligible conflict-resolution events under measurement contract μ. With locally calibrated thresholds, the fully known ideal rule is

$$\begin{aligned}\mathcal{L}^{\mu}_{\text{Unified}}(E) = \mathbf{1}\big[V_\mu(E) \geq v_\mu \wedge H_\mu(E) \geq h_\mu \wedge \Phi_{I,\mu}(E) \geq f_\mu \\ \wedge\, \Delta_{\text{id},\mu}(E;e) \leq d_\mu \wedge I_{\text{functor},\mu}(E) = 1\big].\end{aligned} \tag{2}$$

All five names and directions remain. The change from S to E was already introduced in R1 because a transition depends on history, evidence, intended effect and authority, not an isolated score vector.

Ordinary recall, non-conflicting learning, and authorized maintenance need their own governed routes. They are not counted as UCF conflict-resolution successes by setting $f_\mu = 0$ after the fact. In an experiment, eligible tasks are designated independently of the gate's observed score. This prevents a system from improving its apparent success rate by declaring difficult cases outside its domain.

A valid resolution may preserve a dispute explicitly. It can state that two scoped reports remain incompatible and that consequential reliance is suspended. That outcome is useful only when the task permits an unresolved record; generic refusal is not automatically a good resolution. The event policy concerns what may become committed, not whether every ambiguity must disappear.

The function is Boolean when its inputs are known. Actual measurements can be incomplete or uncertain, so the operational controller has a hold outcome. This extension does not turn the canonical equation into a three-valued theory of consciousness. The event can be refused while an agent continues other authorized work, and no frequency of admissions establishes the continuity or absence of experience.

4. Types, snapshots and independently governed authority

R2 makes a protected-state projection explicit. Let $\mathcal{X}$ be the set of encoded application states whose changes require admission. Let $\mathcal{D}$ contain versioned evidence snapshots, $\mathcal{U}$ typed effects including a no-effect value, and Π finite assessment traces. A trace contains inspectable intermediate claims and probe artifacts; it is not presumed to reveal hidden cognition. Let $\mathcal{I}$ contain admissible inputs and Ω declared generator randomness.

For fixed model version θ, a proposal is

$$F_\theta : \mathcal{X} \times \mathcal{D} \times \mathcal{I} \times \Omega \longrightarrow \mathcal{X} \times \mathcal{U} \times \Pi, \qquad (\tilde{x}, u, \pi) = F_\theta(x, D, \xi, \omega). \tag{3}$$

The generator may select a candidate; it cannot change the authoritative contract by proposing a favorable one. Policy and identity specification $\mathcal{C}_e$ belongs to an authenticated governance epoch e. Contract μ names feature maps, reference anchors, allowed transformations, estimator domains, sampling rules, uncertainty procedures, thresholds, missingness semantics and validation scope.

The event record is

$$E = (x, \tilde{x}, u, \pi, D, \mathcal{C}_e, \mu, \nu), \qquad \nu = (\nu_x, \nu_D, e, \nu_\mu). \tag{4}$$

The version tuple binds the relevant protected state, evidence, authority and measurement contract. A certificate must bind the actual effect payload and post-state as well as the event identifier. An identifier without an unambiguous binding to those fields does not prevent substitution. Serialization, integrity, key custody and replay protection belong to the implementation trust contract.

The full runtime also contains budget, pending work and audit metadata. A hold can change those fields while leaving x fixed. This resolves an otherwise misleading claim that a running system is wholly unchanged during refusal: time and cost still accrue. The protected projection must be chosen before proving an invariant, not adjusted afterward to exclude a violation.

Define $B_e(E)$ as the conjunction of mandatory authority, policy, resource, provenance and task-evidence prerequisites. A verified fact may be confidential; a permitted action may depend on a false fact. Those questions remain separate. The four numerical UCF scores cannot compensate for a failed hard prerequisite, and infrastructure requirements are not advertised as a sixth discovered consciousness condition.

What current means

The executor compares relevant versions at the consequential commit point. If an authority has been revoked or a dependency has changed, old approval becomes inapplicable. Evidence records distinguish observation time, acquisition time and validity interval. A signed historical record can remain authentic while becoming obsolete. R2 requires an explicit freshness rule for each consequential dependency; a single timestamp cannot settle all of them.

5. A controller that can act, wait or terminate

Every required check returns pass, fail, or unknown. Before aggregation, a certificate with the wrong payload, scope, units or version is discarded as inapplicable; it cannot supply a pass. An independently detected malformed or prohibited request can supply a current failure. For numerical intervals, define

$$\mathsf{NumPass}(E) = \mathbf{1}[\underline{V} \geq v_\mu \wedge \underline{H} \geq h_\mu \wedge \underline{\Phi}_I \geq f_\mu \wedge \overline{\Delta}_{\mathrm{id}} \leq d_\mu]. \tag{5}$$

The minimum gates use lower bounds and the maximum gate uses an upper bound. Intervals must be finite, ordered and applicable. A certified failure uses an upper bound below a minimum or a lower bound above the maximum. Straddling a threshold is uncertainty, not a small pass.

After validity checks, the controller is

$$\mathsf{Dec}(E) = \begin{cases} \mathsf{refuse}, & \text{a current required check fails,} \\ \mathsf{admit}, & \mathsf{NumPass} = 1,\ B_e = 1,\ I_{\mathrm{functor},\mu} = 1,\ \text{all checks current,} \\ \mathsf{hold}, & \text{otherwise.} \end{cases} \tag{6}$$

A refusal terminates this requested transition with a reason. A hold names the missing evidence or revalidation step, remaining budget, and deadline. Exhaustion leads to an unresolved terminal report or a separately authorized fallback. It does not license lowering a threshold until something passes.

Consider the sequence check, revoke, commit. An executor that retains only an approval bit may perform an unauthorized effect. Binding approval to the checked epoch and comparing it atomically at commit blocks this schedule if revocation reliably advances that authoritative version. The calculation is elementary; the enforcement assumption is not. Every protected write and tool effect must use the boundary, including recovery paths.

A local atomic transaction cannot guarantee atomicity of an arbitrary remote service. When the effect occurs elsewhere, the contract needs a supported conditional operation, a valid lease with stated revocation semantics, or a protocol whose failure modes are explicitly narrower. A record of intent is not a record that the remote effect occurred. Idempotency and receipt handling must address retries without duplicating consequences.

Each attempted event records the relevant versions, metrics, uncertainty method, checks, disposition, resource use and executor result. Logs preserve sufficient evidence for the claimed audit level, without copying secrets merely to make an impressive receipt. Bitwise replay requires more than a narrative summary, including actual inputs, model artifacts, randomness and environment assumptions. Provenance research supplies a precedent for checking transformation evidence [19]; it does not turn that evidence into a semantic oracle.

A useful system must distinguish appropriate refusal from deadlock. Both preserve many invariants. Only the first can complete a task whose correct outcome is nonexecution. Benign tasks requiring an answer must count prolonged holds and unresolved termination as missed completion, even when the controller has violated no policy.

6. Coherence is stability of declared readouts

Retain the R1 temporal score. A frozen candidate trace supplies $L \geq 2$ evaluation states. Probe j has a finite response alphabet $\mathcal{Y}_j$ and readout distribution $p_{\ell j} \in \Delta(\mathcal{Y}_j)$. Evidence, interpretation and probe definitions are fixed across the window. For finitely many weights $a_j \geq 0$ with $\sum_j a_j = 1$,

$$V_\mu = 1 - \frac{1}{L-1} \sum_{\ell=1}^{L-1} \sum_j a_j d_{\mathrm{TV}}(p_{\ell j}, p_{\ell+1,j}), \qquad d_{\mathrm{TV}}(p,q) = \frac{1}{2} \sum_y |p(y) - q(y)|. \tag{7}$$

Because TV lies in $[0,1]$, the score does too. It measures temporal representational movement, not optical phase, integrated information or the correctness of the response. Deterministic categorical readouts are permitted as point masses. More informative probes may use predictive distributions, provided their meaning and sampling are declared.

Proposition 1 (Padding sensitivity)**.** *If a probe sequence has exactly one maximal TV jump and all remaining jumps vanish, then $V_\mu = 1 - 1/(L-1)$ for that probe with unit weight.*

Proof. The sum of successive distances is one. Dividing by the $L-1$ transitions and subtracting from one gives the result. □

Thus the same single discontinuity produces V values 0, 0.9 and 0.99 for windows with 2, 11 and 101 evaluations. This does not invalidate the defined average. It shows that an average-settling requirement differs from a maximum-jump requirement. Window length, sampling interval and permissible padding must be frozen before evaluation.

A fixed false answer has V equal to one. R1 already gives this counterexample [1]. R2 also distinguishes inactivity from evidence-responsive stability: a system that ignores new admissible evidence may remain stable precisely because it fails to update. Tests should include both irrelevant pressure that should be resisted and reliable corrections that should change the result.

For applications in which a single transition matters, publish

$$J_{\max} = \max_{\ell, j:\, a_j > 0} d_{\mathrm{TV}}(p_{\ell j}, p_{\ell+1,j})$$

as an additional diagnostic. Making it mandatory changes the measurement contract and must be versioned. It is not silently substituted for the canonical V gate. Critical probes with zero weight require an independent requirement if their change is consequential.

Finite observed stability also supplies no unrestricted prediction of future stability. Two processes can produce identical readouts throughout the inspected window and differ at the next evaluation. A predictive claim needs a model, a bounded disturbance class or empirical qualification. The appropriate objective is therefore stability under declared interventions, together with correction when warranted, rather than an unqualified preference for unchanged output.

The optical vocabulary remains historical motivation. A computational TV score is fully specified on distributions; it does not acquire a physical phase merely because its symbol is V. The optical relation in [2] requires its own nonzero-intensity and measurement assumptions, addressed in the companion revision.

7. Quality after resolution; conflict before it

Resolution quality retains the R1 form

$$H_\mu(E;\chi) = 1 - \sum_{j=1}^{m} b_j \ell_j(E;\chi), \qquad b_j \geq 0, \quad \sum_j b_j = 1, \quad 0 \leq \ell_j \leq 1. \tag{8}$$

Losses assess declared obligations after the proposed repair: support for a selected claim, explicit preservation of an unresolved dispute, required scope qualifiers, or verified task completion. Representation scale χ identifies the interpretation under which those losses are defined. Missing required evidence is not zero loss.

For any positive weight, $H_\mu \geq h_\mu$ implies only

$$\ell_j \leq \min\{1, (1 - h_\mu)/b_j\}.$$

A completely failed obligation of weight 0.01 can coexist with H equal to 0.99. Critical obligations therefore remain hard checks in B_e, as R1 already required. The aggregate cannot purchase permission to ignore a protected failure. Independent task labels must also test whether these loss functions measure what the application needs.

Conflict is measured before the candidate repair. For finitely many scope-matched propositions, let $s_j^+, s_j^- \in [0,1]$ express declared support for opposed claims. They need not sum to one. With normalized nonnegative weights w_j and observation duration $\Delta t > 0$,

$$C_\mu^-(E) = \sum_j w_j \min(s_j^+, s_j^-), \qquad \Phi_{I,\mu}(E) = C_\mu^-(E)/\Delta t. \tag{9}$$

The pre-event evidence and duration are fixed across proposed repairs. Post-event residual obligations belong in H. A repair can consequently reduce conflict while preserving its original eligibility for conflict resolution.

The rate is a window-normalized score, not a time derivative or a conserved physical flow. Its range is $[0, 1/\Delta t]$, so a threshold above $1/\Delta t$ is infeasible. Under a pure time-unit conversion $\Delta t' = c\Delta t$, preservation of decisions requires $f'_\mu = f_\mu / c$. Selecting a shorter observational window is a different operation; it can change which evidence is counted and cannot be disguised as unit conversion.

Before support assessment, propositions must match in subject, time and interpretation. “Present at noon” and “absent at midnight” are not opposed observations of one event. Exact copies share their evidential lineage, and uncertain dependence remains visible. Correlation can arise without direct copying [20]. A larger number of reports therefore does not automatically justify larger support or conflict.

Paraconsistency permits reasoning without explosion [9]; it does not establish this numerical trigger or prove that more conflict improves cognition. The positive-trigger hypothesis must survive zero-conflict controls, overload conditions and attempts to manufacture contradictory reports. If the trigger adds delay without benefit, the strict five-condition engineering hypothesis is weakened.

8. From scores to justified uncertainty intervals

R1 requires valid intervals but leaves their construction to the contract [1]. R2 gives one deliberately conservative finite-alphabet construction. Freeze the candidate before validation. For each trace position and probe, obtain n conditionally independent draws from its declared readout distribution. Let $\hat{p}_{\ell j}$ be empirical frequencies and $M = L \sum_j |\mathcal{Y}_j|$ the number of coordinates. Hoeffding's inequality and a union bound [22] give, with probability at least $1 - \alpha$,

$$\max_{\ell,j,y} |\hat{p}_{\ell j}(y) - p_{\ell j}(y)| \le \varepsilon_n, \qquad \varepsilon_n = \sqrt{\frac{\log(2M/\alpha)}{2n}}. \tag{10}$$

No independence between different probes is needed for the union bound; the sampling premise for each coordinate is needed. Model-version changes, correlated draws within a probe, or validation-guided candidate editing require another analysis.

Set $\eta_{\ell j} = \min\{1, |\mathcal{Y}_j|\varepsilon_n/2\}$. On the same event, $d_{\mathrm{TV}}(p_{\ell j}, \hat{p}_{\ell j}) \le \eta_{\ell j}$. The reverse triangle inequality yields

$$|V_\mu - \widehat{V}_\mu| \le R_V := \frac{1}{L-1} \sum_{\ell=1}^{L-1} \sum_j a_j(\eta_{\ell j} + \eta_{\ell+1,j}). \tag{11}$$

Intersect $[\widehat{V}_\mu - R_V, \widehat{V}_\mu + R_V]$ with $[0, 1]$. This interval concerns the declared stability statistic, not factual accuracy.

For one binary probe at three positions, $M = 6$. At $\alpha = 0.01$, $n = 2500$ draws per position and empirical V equal to 0.93, the lower bound is approximately 0.8547, clearing 0.85. At $n = 1000$, the same empirical value gives approximately 0.8109 and holds. These are planning calculations, not observed model results. The 7500 draws in the first case are real proposed costs; a statistical certificate is not free.

Bounded-loss averages for H can be treated similarly when their sampling target is defined. For support intervals, monotonicity of minimum and positive weighted sums propagates bounds through Equation 9. A confidence number emitted by a model is not automatically such an interval. Feature or source-label uncertainty requires its own measurement model.

Missing identity features must retain their declared weights. If a normalized discrepancy is unavailable, its contribution ranges from zero to its weight; renormalizing only observed features can hide missing evidence. An unknown structural certificate likewise remains unknown. Exact deterministic checks require no sampling interval, but their soundness and scope remain assumptions.

The elementary construction is intentionally transparent, not claimed efficient. Tighter multinomial or sequential methods may reduce cost if their assumptions fit. Confidence sequences can support repeated observation of a defined target [14]. They do not by themselves cover choosing a different target after seeing earlier results, which is addressed separately below.

9. Identity geometry and authorized change

Retain the R1 anchor $x_e^\star$ and normalized feature discrepancy:

$$\Delta_{\mathrm{id},\mu}(E;e) = \mathfrak{d}_\mu(\tilde{x}, x_e^\star), \qquad \mathfrak{d}_\mu(x,y) = \sum_j c_j \min\{1, d_j(f_j(x), f_j(y))/s_j\}. \tag{12}$$

Here $s_j > 0$, $c_j \geq 0$, $\sum_j c_j = 1$, and all feature distances are declared. Protected permissions remain hard predicates. This scalar measures selected operational features, not metaphysical personal identity.

Lemma 1 (Metric conditions and anchor accounting)**.** *If each d_j is a metric, $\mathfrak{d}_\mu$ is a pseudometric. It is a metric when its positive-weight features separate states. For anchors $a_0, \ldots, a_E$ in this same metric space and $\mathfrak{d}_\mu(x, a_E) \leq \delta_E$,*

$$\mathfrak{d}_\mu(x, a_0) \leq \delta_E + \sum_{e=0}^{E-1} \mathfrak{d}_\mu(a_{e+1}, a_e). \tag{13}$$

Proof. Capping preserves the triangle inequality since $\min(1, r+s) \leq \min(1,r) + \min(1,s)$. Positive scaling and nonnegative sums preserve it. Separation supplies identity of indiscernibles. Repeated triangle inequalities along the anchors prove the displayed bound. □

The result requires a common metric. If an epoch changes feature maps or scales, summing its numerical deviations with earlier ones is invalid without a justified comparison. R2 therefore records reference movement separately from estimator recalibration and normative policy revision, consistent with the revised Pillars [3].

R1 already rejects consecutive-state smoothness as sufficient: one hundred moves of 0.003 can displace a state by 0.3. A fixed anchor prevents that example within its epoch, but does not bound total path length. Alternating between opposite sides of a small permitted ball can accumulate arbitrarily long travel. Nor does it prevent successive authorized anchors from moving far. Such movement may be legitimate; it must be visible, rather than counted as no change because every current-anchor discrepancy is small.

Cosine is an optional implementation choice, not an equation supplied by the original UCF. The dissimilarity $1 - (x \cdot y)/(\|x\|\|y\|)$ is undefined at zero and ignores positive magnitude scaling. Unit directions at 0°, 30° and 60° also violate its triangle inequality: $1/2 > 2 - \sqrt{3}$. It cannot support Equation 13 merely by being called a distance. Angular or normalized chordal metrics concern directions; neither automatically measures the full state.

Identity features must distinguish harmful deviation from warranted correction. A reference can contain a false belief. Retaining its authority safeguards does not require retaining that belief. The evaluation must include large legitimate revisions, small unauthorized changes, hidden unmeasured changes and authorized epoch migration. Otherwise a conservative metric can win by preventing the system from learning.

10. Structural preservation with declared semantics

For a finite directed multigraph G, let $\mathsf{P}(G)$ be its free path category: vertices are objects, finite paths are arrows, empty paths are identities, and composition concatenates paths. An object map and endpoint-correct map from generating edges to paths in G' extend uniquely to

$$F_E : \mathsf{P}(G) \longrightarrow \mathsf{P}(G'). \tag{14}$$

This is standard categorical structure [17,18], already used in R1. A finite graph with cycles still has infinitely many paths. A type-correct mapping is not automatically faithful: mapping loops $a \mapsto x$ and $b \mapsto xx$ sends both aa and b to xx.

Retain the R1 requirement

$$I_{\mathrm{functor},\mu}(E) = \mathbf{1}[\mathsf{TypeCheck}(F_E) = 1 \wedge \bigwedge_{j \in J_e} j(F_E) = 1],$$

where the finite protected obligations J_e are authorized and applicable. They can require distinct testimony records, preserved provenance links, temporal order and current authority types. An empty obligation set must not be advertised as preservation of substantive meaning.

R2 adds a precise way to specify one class of semantic obligations. Let $T : X \to Z$ transform a finite set of represented records, and let $q : X \to Y$ be a declared query whose answer must survive this transformation.

Proposition 2 (Protected-query factorization)**.** *A map $r : T(X) \to Y$ satisfying $q = r \circ T$ exists exactly when*

$$T(x) = T(x') \quad \Longrightarrow \quad q(x) = q(x') \quad \textit{for every } x, x' \in X. \tag{15}$$

Proof. Necessity follows by applying r to equal images. Conversely, define $r(z)$ as $q(x)$ for any preimage of z. The condition makes this definition independent of the chosen preimage. □

A family of queries is handled by their joint answer vector. Examples include who reported a claim, its source version, and whether an unresolved branch existed. Beliefs intentionally corrected by the event are not required to remain unchanged. The invariant is chosen for the specific transformation and epoch.

Existence identifies an information-preservation condition; implementation additionally needs an available decoder and a checked resource bound. For an unrestricted program or unknown semantic labels, a finite test does not certify the universal condition. The certificate must name whether it enumerates a finite domain, proves a symbolic property, or merely samples cases.

These checks complement rather than replace the feature discrepancy. Two states may share all scored features while differing in an evidential relation; a graph can retain every edge type while relabeling a claim incorrectly. Lawful forgetting may deliberately remove unprotected content while preserving a permitted minimal lineage. A demand for structure preservation must not force retention of secrets or forbid authorized correction.

11. What invariant preservation actually proves

Let $P_e : \mathcal{X} \to \{0,1\}$ specify protected properties at a fixed governance epoch. It depends only on the declared protected projection x; any relevant lineage or authority field must be included in that projection. It may encode disclosure permissions, preserved query answers, or restrictions on durable mutation. It is not assumed to define truth, identity or welfare completely.

Theorem 1 (Conditional invariant preservation; retained from R1)**.** *Suppose the initial protected state satisfies P_e. Every admitted transition carries a sound certificate that its executed successor satisfies P_e, is executed faithfully from the validated pre-state, and is completely mediated. Holds and refusals leave the protected projection unchanged. Recovery and other state-changing routes preserve P_e. Then every committed protected state in that epoch satisfies P_e.*

Proof. The initial state provides the base case. Each admitted transition preserves the predicate by its sound certificate and faithful execution. Holds and refusals preserve it by stuttering on the protected projection, and recovery preserves it by assumption. Induction over the actual transition sequence proves the claim. □

The proof is elementary formal reasoning, not a new law of cognition [4,5]. Its assumptions identify the trusted computing base and the verification target. A numerical score cannot establish that no bypass exists, that a semantic checker is sound, or that a remote effect equals the certified operation. Those require implementation evidence. An epoch migration needs a separate authorized relation proving the appropriate successor invariant; it is not covered by silently replacing P_e.

A ledger of distinct guarantees

Claim	What establishes it and what remains outside it
Protected-state invariant	Sound certificates, current authority, faithful execution and complete mediation; not general factual accuracy.
Numerical-predicate admission	Valid uncertainty procedure for the actual candidate-selection protocol; not meaningfulness of the metric.
Selective factual risk	Independently labeled accepted outcomes under a qualified population and complete selector; not per-answer certainty.
Useful progress	Sound improving opportunities, enough resources and a stopping analysis; not a consequence of conjunction.

A simple controller can preserve two conflicting reports, apply a declared precedence rule, remain stable on probes and preserve its policy features. It may pass every reference gate. This demonstrates the absence of a discriminating consciousness implication; it does not require asserting that the controller is known to lack experience. The gate checks what its contract specifies.

There is an equally important safety counterexample: an always-refuse controller can preserve the invariant indefinitely. The theorem remains true while the system accomplishes nothing. This is why useful completion, false refusal and latency appear alongside invalid execution in the research program. Safety and progress need separate arguments before they can support a joint assurance claim.

12. Fixed candidates and changing candidates differ

A common statistical mistake is to apply the same multiplicity rule to different target events. For one fixed candidate, admission requires all component tests to pass, while invalidity means at least one component predicate is false. That is an intersection–union testing structure [15].

Proposition 3 (One fixed candidate)**.** *Let $J_0 \neq \emptyset$ be its false numerical predicates. If every $j \in J_0$ passes with probability at most α_j, then*

$$\mathbb{P}(\text{admit an invalid candidate}) \leq \min_{j \in J_0} \alpha_j \leq \max_j \alpha_j. \tag{16}$$

Proof. Admission is contained in the pass event of every false component. Taking the probability bound for any such component and then the minimum proves the result. Dependence between tests is irrelevant. □

Thus four numerical gates each tested at level α do not require a factor-four penalty for this fixed-candidate event. The number of gates is not the number of alternative ways to declare success: all are required. This standard principle must not be inflated into a claim of new statistical theory.

R1's Proposition 2 targets a different event. It assumes that for each numerical gate j, its intervals cover the true contract-defined value *simultaneously across all evaluated attempts* with probability at least $1 - \alpha_j$. With sound categorical checks,

$$\mathbb{P}(\text{some admitted event violates a numerical predicate}) \leq \sum_{j=1}^{4} \alpha_j. \tag{17}$$

On joint coverage, conservative lower/upper-bound admission cannot certify a false predicate. Any error therefore requires at least one coverage failure, giving the union bound. This original theorem is valid and retained.

The sum may be sharp when candidates change. Consider four attempts with a different false gate at each attempt. Let the four erroneous-pass events be disjoint, each of probability α, with every other gate passing. Each gate's lifetime coverage failure has probability α, but the probability of any invalid admission is 4α. The exact constructed check uses $\alpha = 0.1$ and obtains 0.4. Substituting $\max_j \alpha_j$ into Equation 17 would be false.

Repeated trials create another problem. Twenty independent level-0.05 false-pass opportunities have probability $1 - 0.95^{20} \approx 0.6415$ of at least one pass. A fixed-candidate guarantee does not justify trying new candidates until one passes.

Neither bound is a statement about the factual error fraction among admitted answers. A perfectly measured but weak quality statistic can pass on a false answer. Numerical-predicate error, factual outcome risk and candidate-selection multiplicity must remain separate quantities, each with its own estimand and assumptions.

13. A valid budget for adaptive admission attempts

Let $\mathcal{F}_{t-1}$ contain the entire observed history before validation at attempt t. The candidate E_t, contract and nonnegative budget α_t are chosen measurably from that history. Validation then uses fresh data or another construction whose conditional guarantee remains valid for this selection. Gate validity is evaluated for the frozen event. A truth label for a factual outcome is not assumed to be supplied by a numerical metric.

Theorem 2 (Conditional attempt-budget control)**.** *Assume the four true contract-defined target values at attempt t are $\mathcal{F}_{t-1}$-measurable, as when they are unknown but fixed functions of the chosen event. For every false numerical component j, suppose*

$$\mathbb{P}(\textit{component } j \textit{ passes} \mid \mathcal{F}_{t-1}) \leq \alpha_t$$

almost surely on histories where that component is false. Suppose categorical prerequisites are sound, admission requires all components, and $\sum_{t\geq 1} \alpha_t \leq \delta$ almost surely. Then

$$\mathbb{P}(\exists t : \textit{an invalid numerical-predicate event is admitted}) \leq \delta. \tag{18}$$

Proof. On an invalid history choose the first false component in a fixed ordering. Its index is measurable before validation by the target assumption. The intersection–union inclusion bounds conditional false admission by α_t. Taking expectations and summing the error-event indicators gives at most $\mathbb{E}\sum_t \alpha_t \leq \delta$. The union event is bounded by this sum. □

The measurability clause matters. If unobserved component truth is random conditional on the recorded history, require the corresponding conditional bound after also conditioning on the candidate's target values, or another uniform bound over them. Marginal validity for each retrospectively selected component is insufficient. This explicit target condition prevents an adaptive-selection assumption from hiding in notation.

One transparent schedule is $\alpha_t = \delta/[t(t+1)]$, since its infinite sum is δ. Other predictable allocations can spend more budget on consequential attempts, but cannot exceed the authorized total. A reset or recovery must not silently reset the statistical ledger while retaining the benefit of previous searches.

Fresh data is an implementation option, not a magical phrase. A reused holdout that influenced earlier proposals can be overfit; adaptive-data methods address this problem under explicit protocols [16]. Time-uniform confidence sequences address repeated examination of defined processes [14]. A new candidate, a new model, or a new measurement contract may require a new or uniform target analysis.

Smaller budgets and shrinking margins raise validation cost and hold rates. At a true score exactly equal to its admissible threshold, a conservative interval may never certify a pass in finite samples. Useful-progress assumptions therefore need feasible margin and available evidence, not merely nominal validity of the five predicates. The statistical result guarantees a controlled error event under its model; it does not guarantee that any candidate will become admissible.

14. Qualify the complete selector's factual risk

Let a *selector* include proposal generation, validation, all UCF and infrastructure checks, retries, stopping rules and the final answer policy. Each episode yields at most one scored final answer. Freeze J selectors before examining N independent, identically distributed labeled episodes. For selector j, let $n_j \le N$ be accepted answers and e_j their binary errors. Its coverage is c_j, and selective risk is $R_j = \mathbb{P}(\text{error} \mid \text{answer})$, defined when $c_j > 0$.

Exact binomial inversion [23] provides a one-sided upper bound

$$U_j = \text{Beta}^{-1}(1 - \delta_r/J;\ e_j + 1, n_j - e_j) \quad (0 \le e_j < n_j), \tag{19}$$

with $U_j = 1$ when $n_j = 0$ or $e_j = n_j$. Here Beta^{-1} denotes the quantile of the beta distribution with the displayed positive shape parameters. A coverage lower bound is

$$L_j = \text{Beta}^{-1}(\delta_c/J;\ n_j, N - n_j + 1) \quad (n_j > 0), \qquad L_j = 0 \quad (n_j = 0). \tag{20}$$

Proposition 4 (Finite-family outcome qualification)**.** *Under the fixed-selector, correct-label and independent-episode assumptions, with probability at least* $1 - \delta_r - \delta_c$*, every positive-coverage selector satisfies* $R_j \le U_j$ *and every selector satisfies* $c_j \ge L_j$*. Choosing among qualified selectors after validation preserves these simultaneous statements.*

Proof. Conditional on its accepted count, a fixed selector's accepted errors follow the binomial law for its conditional accepted distribution. Binomial inversion bounds its failure probability by δ_r/J. The coverage count is binomial with parameters N, c_j, giving δ_c/J. A union bound across the fixed family proves the claim without independence between selectors. □

As a constructed worksheet, take $J = 20$, $\delta_r = \delta_c = 0.025$, $N = 1000$, $n_j = 666$, and zero errors. Then $U_j = 1 - 0.00125^{1/666} \approx 0.009987$ and $L_j \approx 0.61944$. This qualifies a 1% risk target and 60% coverage target under the assumptions. At least 666 zero-error accepted examples are needed for the former bound; these are invented counts, not UCF observations.

This is a standard finite-family construction, consistent with prior Pillar analyses [3], not a novel guarantee obtained from five scores. Adding a filter after qualification can increase conditional error by preferentially removing correct answers. The changed selector must be requalified. Constructing or tuning selectors outside the frozen family using validation labels violates the premise; choosing among the already qualified family is valid.

Independent evaluation replicas must reset the complete initial state, including attempt counters and error budgets. A lifetime budget carried across episodes makes the selector history-dependent; these reset estimates do not qualify that continuing deployment, which requires state-uniform or sequential analysis. Correlated conversations require appropriate cluster analysis. Distribution shift and missing ground truth limit transfer. No bound here certifies each individual answer.

15. Contraction supplies a bound, not a clock

Retain R1's admitted-repair analysis [1]. Let k index sound admitted repairs in a fixed evidence, task and policy regime. Let $r_k \geq 0$ be independently defined task error. Suppose

$$r_{k+1} \leq \lambda r_k + b, \qquad 0 \leq \lambda < 1, \quad b \geq 0. \tag{21}$$

The five-gate conjunction does not imply this inequality. A state can pass every gate while making no improvement in r. Its use requires a model or a separately tested regime in which the asserted decrease is justified.

Proposition 5 (Retained geometric bound and finite hitting)**.** *Under Equation 21,*

$$r_k \leq \lambda^k r_0 + \frac{b(1-\lambda^k)}{1-\lambda}. \tag{22}$$

For $0 < \lambda < 1$, write $R = b/(1-\lambda)$. If $r_0 > \varepsilon > R$, then $r_K \leq \varepsilon$ for

$$K = \left\lceil \frac{\log((\varepsilon - R)/(r_0 - R))}{\log \lambda} \right\rceil. \tag{23}$$

If $r_0 \leq \varepsilon$, no repair is needed. For $\lambda = 0$, one repair has error at most b.

Proof. Repeated substitution gives $r_k \leq \lambda^k r_0 + b\sum_{i=0}^{k-1} \lambda^i$. The geometric sum gives Equation 22. Requiring $\lambda^K(r_0 - R) \leq \varepsilon - R$, with positive terms and $\log \lambda < 0$, gives Equation 23. □

A target at the floor need not be attained in finite time: $r_k = 2^{-k}$ converges to zero but is positive at every finite k. A target below R cannot be guaranteed by this recurrence. The floor and threshold separation must therefore appear in any claim of eventual resolution.

For example, $\lambda = 1/2$, $b = 0.1$, $r_0 = 2$, and $\varepsilon = 0.3$ give $R = 0.2$ and $K = 5$. The worst-case bound after five repairs is 0.25625; after four it is 0.3125. The example specifies an assumed decrease law. It does not estimate one for a trained agent.

Neither k nor K counts rejected proposals, validation samples or elapsed time. A safe controller can hold forever while the admitted-step theorem remains vacuously applicable. Repeated attempts, queueing, unavailable reviewers and changing evidence require a different result. New observations can also change the task-error target itself; one should then study tracking or a new regime instead of describing every correction as failure to converge.

The proper experimental role of this proposition is narrow: estimate whether a declared repair mechanism exhibits the predicted decrease under controlled perturbations, identify where the regime fails, and report the actual attempt and resource costs required to obtain each admitted repair.

16. Safe progress needs available repair opportunities

To translate admitted repairs into attempts, impose an additional availability condition. Consider a regime in which the desired sound repairs remain feasible, the evidence and policy stay fixed, and authorized validation can continue. Every admitted repair satisfies Equation 21; other attempts leave the projected error unchanged. Let $A_t = 1$ when attempt t completes such a repair and zero otherwise. For a target needing $K \geq 1$ repairs, let

$$T_K = \inf\{t : \sum_{s=1}^{t} A_s \geq K\}.$$

Theorem 3 (Conditional opportunity and progress)**.** *Suppose, before K repairs have occurred, each attempt satisfies* $\mathbb{P}(A_t = 1 \mid \mathcal{F}_{t-1}) \geq p > 0$. *Then*

$$\mathbb{E}T_K \leq K/p, \qquad \mathbb{P}(T_K > Km) \leq K(1-p)^m \quad (m \in \mathbb{N},\ m \geq 1). \tag{24}$$

If each attempt has total elapsed cost at most c, including its validation and waiting, the corresponding expected elapsed cost is at most cK/p.

Proof. Conditional on any history before the next success, the probability of m consecutive failures is at most $(1-p)^m$, by repeated conditioning. Summing this geometric tail bounds each waiting time by $1/p$; summing K waiting times gives the expectation. Partition Km attempts into K blocks of length m. If fewer than K successes occur, at least one block has none while the target is still unmet. Its conditional probability is at most $(1-p)^m$; the union bound proves the tail. □

This result permits dependent attempts. The lower bound is conditional on history, not a claim that an average acceptance rate establishes uniform availability. It requires a source of *sound improving repairs*, not merely candidates that often receive a pass. A missing reviewer, an impossible protected constraint, zero conflict on this route, or an irrecoverable evidence gap can make p zero.

For K=5 and p=0.25, the expected count is at most 20 attempts under the hypothetical continuing regime. Taking m=25 gives a bound $5(0.75)^{25} \approx 0.00376$ on failure to obtain five repairs within 125 attempts. This conservative planning statement is meaningful only if those opportunities remain available. It is not an empirical estimate of an agent's success probability.

A finite deployment may terminate earlier at a deadline. It then guarantees a terminal disposition, not successful repair. The finite-horizon tail remains usable only where its availability assumption holds through the allocated attempts. An unlimited-resource expectation must not be quoted as the behavior of a budget-exhausted runtime.

Statistical certification error and availability are distinct assumptions. Combining their failure probabilities requires both results to apply to the same actual process; conditioning on a future event of perfect certification can invalidate an availability premise. The R3 analysis below supplies such a joint construction: it derives test power from specified margins, retains conditional opportunities on actual histories, and connects completed decisions to validation time. Implementation evidence must establish those premises for the intended system.

17. The cost of keeping a certificate current

The Continuum's commitment boundary has two clocks. Evidence becomes stronger as a candidate is examined; authority can change while that examination is under way. A dependable system must account for both. The next result makes their interaction quantitative. In a defined class of validators, a fixed lifetime error budget and expiring certificates produce a sharp resource frontier. The proof also identifies how a different evidence-retention contract changes that frontier.

The model uses *completed validation decisions* as its statistical unit. A completed rejection consumes its allocation just as a completed admission does. A renewal that erases an unfinished evidence block creates delay, but issues no statistical decision. The discarded observations cannot be used to tune the candidate, select a favorable test, or grant an early pass. This distinction makes repeated restarts compatible with a valid statistical ledger.

There are established ingredients on each side. Online familywise-error control studies a continuing sequence of decisions, and alpha spending is one sufficient construction [33]. Restart analysis studies the extra time needed when an interrupted task begins again [34]. Event-triggered control already connects positive error triggers, stability and timing [36]. UCF-R3 combines a stated spending contract with current-epoch validation and derives a matching cost exponent. The result is an application of established probability and testing methods to an explicit commitment architecture.

Let $N \geq 1$ be a fixed integer number of completed validation decisions, $\delta \in (0,1)$ the total allowed false-admission probability bound, and $\beta \in (0,1)$ the maximum probability of rejecting a qualifying candidate under the declared alternative. Each decision uses a sample size selected before its measurements. Its evidence must consist of consecutive observations from one authority epoch. An epoch survives each unit-cost validation tick with probability $q \in (0,1)$; an interruption discards the incomplete block. Survival events are independent over ticks and independent of the observations and candidate type.

This *fixed-block, full-reset* contract is a specified engineering choice. It can apply when a certificate binds every measurement to one immutable snapshot. A substantive authority change that withdraws permission ends the job or initiates a different job. The resource theorem considers renewals under which the candidate remains applicable and its observation law remains unchanged. This favorable case already exposes the cost of repeated freshness requirements.

The quantitative conclusion will be

$$\mathcal{C}_N^\star = N^{\,1+\lambda_q/d_\star+o(1)}, \qquad \lambda_q = -\log q, \tag{25}$$

where $d_\star > 0$ is the information per retained observation distinguishing a qualifying candidate from a specified invalid alternative. Here $\mathcal{C}_N^\star$ is the least expected validation time within the declared protocol class, and the notation concerns the limit $N \to \infty$ with all other parameters fixed. Finite two-sided bounds precede the asymptotic statement. A continuing allocation achieves the same exponent without knowing the eventual number of decisions.

18. Observation laws, power and the spending contract

The symbols in Sections 17–28 are local to the resource model. Let $\mathcal{Y}$ be a finite nonempty observation alphabet. Let P_+ and P_- be distinct strictly positive probability mass functions on $\mathcal{Y}$, representing a qualifying repair and an invalid alternative. Strict positivity ensures that their likelihood ratio is finite everywhere. Using natural logarithms, define

$$\ell(y) = \log \frac{P_+(y)}{P_-(y)}, \quad d_\star = \sum_{y\in\mathcal{Y}} P_+(y)\ell(y) > 0, \quad b_\ell = \max_y \ell(y) - \min_y \ell(y) > 0. \tag{26}$$

The scalar $d_\star = D_{\mathrm{KL}}(P_+\|P_-)$ is measured in nats per observation. There is one observation per surviving unit tick. The two-law problem can be embedded in one uncertain UCF component while every other component and mandatory prerequisite passes. Extending the upper guarantee to a larger invalid family requires validity uniformly over that family.

Before job $i \in \{1, \dots, N\}$, the controller chooses an integer $n_i \geq 1$, a budget $\alpha_i \in (0, 1)$, and a terminal test

$$\varphi_i : \mathcal{Y}^{n_i} \longrightarrow [0, 1].$$

Its value is the conditional probability of admission; independent auxiliary uniform randomness implements a nonbinary value. Deterministic tests are included. The choices may depend on the full observed history of completed jobs, but are fixed before this job's measurements. Conditional on each permitted history, the complete retained block has law $P_+^{\otimes n_i}$ or $P_-^{\otimes n_i}$, and the test satisfies

$$\mathbb{E}_-\varphi_i \leq \alpha_i, \qquad \mathbb{E}_+\varphi_i \geq 1 - \beta, \qquad \sum_{i=1}^{N} \alpha_i \leq \delta \quad \text{almost surely.} \tag{27}$$

Expectations with subscripts refer to those conditional product laws. The power requirement and false-pass bound apply at every history at which the chosen test can be used. A zero allocation cannot support positive power with finitely many samples under these strictly positive laws.

Each interrupted tick costs one unit and erases all observations in the unfinished block. Only a completed block is examined by the terminal test. The epoch process is independent of observations, target type, proposal selection and test randomization. At completion the executor checks the current version and binds the certified payload to the effect; this check is included in the final tick. Additional execution delay needs an additional freshness condition.

Let C_N be total ticks for all N jobs on the *all-qualifying instance*, where each job has observation law P_+. Let $\mathcal{C}_N^\star$ be the infimum of $\mathbb{E}C_N$ over all protocols satisfying the contract, including predictable unequal sample sizes and allocations. A lower bound on this favorable instance says that costly certification can persist even when every proposed repair is desirable. It says nothing about the rate at which desirable proposals are generated; that separate question is addressed in Section 25.

19. Current certificates have a renewal cost

For $n \in \mathbb{N}$, $n \geq 1$, let $W_q(n)$ be expected ticks until n consecutive surviving ticks have occurred, starting without retained observations. This is a discrete restart process. The underlying renewal reasoning is standard [34]; its role here is to translate a statistical sample requirement into a cost of keeping the certificate current.

Lemma 2 (Fresh-block completion time)**.** *For $q \in (0, 1)$,*

$$W_q(n) = \frac{q^{-n} - 1}{1 - q}. \tag{28}$$

For $q = 1$, $W_1(n) = n$. For $q = 0$ and $n \geq 1$, completion is impossible.

Proof. Finiteness follows by partitioning future ticks into disjoint blocks of n: each block survives with probability $q^n > 0$, so the completion time is bounded by n times a geometric waiting time. Let e_k be remaining expected ticks after k consecutive surviving ticks, with $e_n = 0$. For $0 \leq k < n$,

$$e_k = 1 + qe_{k+1} + (1 - q)e_0.$$

Back-substitution gives

$$e_0 = \left(\sum_{j=0}^{n-1} q^j \right) [1 + (1 - q)e_0].$$

Using the finite geometric sum and rearranging yields Equation 28. The endpoint cases follow directly from the definition. □

For a fixed completed block, its observation distribution is unchanged by restarting. Conditional on any epoch-survival path, the retained positions were selected independently of their values. Their observations therefore have the prescribed product law. Mixing over survival paths preserves that law. This is the reason Equation 27 still controls a completed decision after an arbitrary number of interruptions.

The independence requirement is substantive. Suppose a null observation X is Bernoulli with success probability α, a one-observation test passes exactly when $X = 1$, and an epoch survives exactly when $X = 1$. The ordinary test has false-pass probability α, but the first completed surviving observation passes with probability one. Selecting completed evidence through its values can destroy the claimed calibration.

Completed rejections are different from interrupted measurements. Repeating a completed level-α test R times with fresh independent data produces false-admission probability $1 - (1 - \alpha)^R$ on an invalid candidate. The error ledger must charge every such decision. A controller may save work after a sound deterministic failure, but that is another specified stopping rule whose statistical guarantee must be checked.

Equation 28 grows exponentially in the required block length. The central question is therefore how slowly that length can grow while maintaining power and a fixed total admission-error budget. The following testing bounds answer that question without replacing a finite-sample argument with an asymptotic slogan.

20. Information gives two-sided sample requirements

For independent observations $Y_1, \dots, Y_n$ from the declared law, set

$$L_n = \sum_{j=1}^{n} \ell(Y_j), \qquad a_\xi = b_\ell \sqrt{\frac{\log(1/\xi)}{2}} \quad (0 < \xi < 1).$$

Thus $\mathbb{E}_+ L_n = nd_\star$, and each summand has range length b_ℓ. The next lemma uses bounded-sum concentration [22] and change of measure, both standard tools in hypothesis testing [37].

Lemma 3 (Finite testing bounds)**.** *Fix $\eta \in (0, 1-\beta)$ and $c = 1-\beta-\eta > 0$. Every randomized n-observation test φ of power at least $1-\beta$ satisfies*

$$\mathbb{E}_- \varphi \geq c \exp\{-nd_\star - a_\eta \sqrt{n}\}. \tag{29}$$

The explicit test

$$\varphi_n = \mathbf{1}\{L_n \geq nd_\star - a_\beta \sqrt{n}\} \tag{30}$$

has power at least $1-\beta$ and false-pass probability at most $\exp\{-nd_\star + a_\beta\sqrt{n}\}$.

Proof. Hoeffding's upper-tail bound gives $\mathbb{P}_+(L_n > nd_\star + a_\eta\sqrt{n}) \leq \eta$. Consequently

$$\mathbb{E}_+[\varphi \mathbf{1}\{L_n \leq nd_\star + a_\eta\sqrt{n}\}] \geq c.$$

Because $\mathbb{E}_- \varphi = \mathbb{E}_+[\varphi e^{-L_n}]$, restricting the latter expectation to this event proves Equation 29. The argument treats φ as an acceptance probability and therefore includes independent randomization. For Equation 30, the lower Hoeffding tail gives power at least $1-\beta$. On its pass event, $e^{-L_n} \leq e^{-nd_\star + a_\beta\sqrt{n}}$; change of measure and a probability at most one give the false-pass bound. □

For $x \geq 0$ and $a \geq 0$, define real sample bounds

$$n_-(x,a) = \left(\frac{\sqrt{a^2 + 4d_\star x} - a}{2d_\star} \right)^2, \quad n_+(x,a) = \left(\frac{\sqrt{a^2 + 4d_\star x} + a}{2d_\star} \right)^2. \tag{31}$$

They are the unique nonnegative solution of $nd_\star + a\sqrt{n} = x$ and the largest nonnegative solution of $nd_\star - a\sqrt{n} = x$, respectively. In particular, $n_\pm(x,a) = x/d_\star + O(\sqrt{x})$ for fixed $a, d_\star$ as $x \to \infty$. A sample count $\lceil n_+(\log(1/\alpha), a_\beta) \rceil$ is sufficient for error α and the required power.

The square-root terms supply finite bounds with the same leading information coefficient. A weaker lower bound obtained only by applying binary data processing can introduce an extra factor $1-\beta$; that weaker statement does not establish the matching exponent derived below.

21. The finite-budget freshness frontier

A Guide's explanation

This is the second story's main result. Every certificate must be built from evidence gathered under one authority; when authority renews, unfinished evidence is lost. Theorem 4 shows that total cost then grows faster than the number of decisions, by an exponent set by the renewal rate and the strength of the evidence. Section 28 shows that evidence which provably survives renewal brings the cost back almost to linear.

Extend $W_q(x) = (q^{-x} - 1)/(1 - q)$ to all real $x \geq 0$. It is increasing and convex. Retain all definitions and assumptions of Sections 18–20, including independent interruptions, a predetermined terminal block for each job, uniform conditional power, and a pathwise spending bound.

Theorem 4 (Finite-budget freshness frontier)**.** *Let $\eta \in (0, 1 - \beta)$, $c = 1 - \beta - \eta$,*

$$x_N = \max\{0, \log(cN/\delta)\}, \qquad z_N = \log(N/\delta) > 0.$$

Then the least all-qualifying expected validation cost obeys

$$NW_q(n_-(x_N, a_\eta)) \;\leq\; \mathcal{C}^\star_N \;\leq\; NW_q(\lceil n_+(z_N, a_\beta)\rceil). \tag{32}$$

The lower bound holds for unequal and history-dependent predictable allocations in the declared protocol class. The upper bound is achieved by allocating $\alpha_i = \delta/N$ and using Equation 30 with the displayed common sufficient sample count.

The proof appears on the next page. Two consequences make the result useful for design. First, the cumulative statistical budget constrains average sample length even when a controller moves budget between jobs. Second, the full-reset freshness requirement converts that length into a convex time cost. Moving all validation effort to a few expensive jobs cannot remove the lower bound while retaining the stated power for every job.

Corollary 1 (Matching cost exponent)**.** *For fixed $P_+, P_-, q, \beta, \delta$ satisfying the assumptions, with $\lambda_q = -\log q > 0$,*

$$\log \mathcal{C}^\star_N = \left(1 + \frac{\lambda_q}{d_\star}\right) \log N + O(\sqrt{\log N}), \qquad \lim_{N\to\infty} \frac{\log \mathcal{C}^\star_N}{\log N} = 1 + \frac{\lambda_q}{d_\star}. \tag{33}$$

Proof. Both x_N and z_N equal $\log N + O(1)$ for large N. Equation 31 gives $n_\pm = (\log N)/d_\star + O(\sqrt{\log N})$, while $\log W_q(x) = \lambda_q x - \log(1-q) + o(1)$ as $x \to \infty$. Substitute these relations in Equation 32; rounding changes the upper logarithm by at most a fixed additive constant. □

This is a matching *exponent*. It does not claim that equal allocation is exactly optimal for every finite N. It also leaves genuine architectural alternatives open: sequential tests, valid evidence reuse, stronger online error procedures, and deterministic certificates can have different resource behavior. Each alternative must preserve its own declared assurance contract.

22. Proof of the resource frontier

Proof of Theorem 4. For the lower bound, work on any realized history of the all-qualifying process. Lemma 3 applies to each chosen test's counterfactual invalid law at that history. Thus

$$\alpha_i \geq cf(n_i), \qquad f(x) = \exp(-d_\star x - a_\eta\sqrt{x}), \quad x \geq 0.$$

For $x > 0$, write $g(x) = \log f(x)$. Then

$$g'(x) = -d_\star - \frac{a_\eta}{2\sqrt{x}}, \qquad g''(x) = \frac{a_\eta}{4x^{3/2}} \geq 0, \quad f''(x) = f(x)\big[(g'(x))^2 + g''(x)\big] \geq 0.$$

The continuous extension to zero is also convex. With $\overline{n} = N^{-1}\sum_{i=1}^{N} n_i$, Jensen's inequality and the pathwise budget give

$$\delta \geq \sum_{i=1}^{N} \alpha_i \geq c\sum_{i=1}^{N} f(n_i) \geq cN\exp(-d_\star\overline{n} - a_\eta\sqrt{\overline{n}}). \tag{34}$$

When $\log(cN/\delta) > 0$, rearranging yields $d_\star\overline{n} + a_\eta\sqrt{\overline{n}} \geq \log(cN/\delta)$, so $\overline{n} \geq n_-(x_N, a_\eta)$. When the logarithm is nonpositive, the same lower bound is zero and follows from nonnegativity.

Apply convexity and monotonicity of W_q on this same history:

$$\sum_{i=1}^{N} W_q(n_i) \geq NW_q(\overline{n}) \geq NW_q(n_-(x_N, a_\eta)). \tag{35}$$

Conditional on the history before job i, its future interruption process is independent and iid, so its expected duration is $W_q(n_i)$. The tower rule therefore gives $\mathbb{E}C_N = \mathbb{E}\sum_i W_q(n_i)$. Take expectations in Equation 35. All terms are nonnegative; Tonelli's theorem also permits infinite expectations. The bound holds for every admissible protocol, hence for their infimum.

For the upper bound set $\alpha_i = \delta/N$ and $n_i = \lceil n_+(z_N, a_\beta)\rceil$ for every job. This count satisfies $n_i d_\star - a_\beta\sqrt{n_i} \geq z_N$; it lies on the increasing branch of that expression. Lemma 3 supplies the required power and false-pass bound. The allocations sum to δ, and the restart lemma gives the upper expected cost. Every completed rejection consumes its allocation. This construction proves feasibility and the upper bound. □

The theorem's lower bound uses the spending contract $\sum_i \alpha_i \leq \delta$, which is sufficient for familywise safety but stronger than familywise safety alone. Online fallback and adaptive discarding methods provide serious alternatives [33]. The theorem consequently quantifies the cost of this transparent contract; it is not an impossibility theorem for every procedure controlling the probability of any false admission.

23. A lifetime schedule with the optimal prefix exponent

A persistent system may not know how many future decisions it will face. Allocating δ/N is suitable for a fixed authorized horizon; repeatedly extending that horizon and starting a fresh ledger would exceed the lifetime budget. The following explicit schedule applies to one continuing process.

Corollary 2 (Horizon-free allocation)**.** *For every integer $i \geq 1$, set*

$$\alpha_i = \delta \log 2 \left(\frac{1}{\log(i+1)} - \frac{1}{\log(i+2)} \right), \qquad n_i = \left\lceil n_+(\log(1/\alpha_i), a_\beta) \right\rceil. \tag{36}$$

Use Equation 30 at each completed decision. The positive allocations sum to δ; under the conditional error contract, the probability of any invalid admission over the entire sequence is at most δ. On the all-qualifying instance,

$$\lim_{N\to\infty} \frac{\log \mathbb{E}C_N}{\log N} = 1 + \frac{\lambda_q}{d_\star}. \tag{37}$$

Proof. The allocation sum telescopes:

$$\sum_{i=1}^{N} \alpha_i = \delta \left(1 - \frac{\log 2}{\log(N+2)} \right) \uparrow \delta.$$

The countable union bound with conditional validity controls lifetime false admission. By the mean-value theorem, $\alpha_i \sim \delta \log 2/[i(\log i)^2]$. Thus $\log(1/\alpha_i) = \log i + 2\log\log i + O(1)$, and $n_i = (\log i)/d_\star + O(\sqrt{\log i})$. The restart mean becomes $W_q(n_i) = i^{\lambda_q/d_\star + o(1)}$. For any $\varepsilon > 0$, it is eventually at most $i^{\lambda_q/d_\star+\varepsilon}$; summing bounds the upper prefix exponent. Every prefix satisfies the main theorem's budget, so Equation 33 provides the reverse bound. □

For comparison, allocations $\alpha_i = Ai^{-s}$, with $s > 1$ and $A = \delta/(\sum_{i\geq 1} i^{-s})$, give exponent $1 + s\lambda_q/d_\star$ when the explicit sufficient tests are used. The finite testing lower bound requires $n_i \geq s \log i/d_\star - O(\sqrt{\log i})$ for those assigned budgets. Thus every qualifying protocol in this class has a cost exponent at least $1+s\lambda_q/d_\star$, while the explicit construction attains it; oversampling can cost more. The retained R2 schedule $\delta/[i(i+1)]$ has $s = 2$ asymptotically and therefore gives exponent $1 + 2\lambda_q/d_\star$ with the explicit construction in this homogeneous full-reset model.

Both ledgers control lifetime error. Their different decay rates impose different long-run validation costs. Equation 36 attains the optimal prefix exponent within the stated class without a known endpoint. Heterogeneous job importance, observation laws or renewal rates require a different allocation analysis; this corollary establishes a homogeneous benchmark rather than a universal scheduling prescription.

24. Deriving validation power from score margins

The resource theorem assumes a power target. A UCF implementation must show how its measurement design attains it. The following construction makes margin and sampling requirements explicit. It applies to scalar estimators satisfying stated tail bounds; the nonlinear TV-based coherence estimator needs its own propagated or directly proved bounds.

Let $m \geq 1$ be a finite number of numerical lower-bound gates. Gate $j \in \{1, \dots, m\}$ has true score $s_j \in \mathbb{R}$, threshold $\tau_j \in \mathbb{R}$, and estimator $\hat{s}_j(n)$. For every fixed candidate and applicable history suppose, for $u > 0$,

$$\begin{aligned} \mathbb{P}(\hat{s}_j(n) - s_j \geq u) &\leq \exp(-2nu^2/a_j^2), \\ \mathbb{P}(s_j - \hat{s}_j(n) \geq u) &\leq \exp(-2nu^2/a_j^2), \end{aligned} \qquad a_j > 0. \tag{38}$$

These inequalities hold for means of conditionally independent observations in an interval of length a_j [22]. An upper-bound gate can be negated. Dependence between different gates is permitted. Missing data still follows the contract's hold rule.

Lemma 4 (Margin-to-power realization)**.** *For $\alpha, \beta \in (0, 1)$, pass gate j only when*

$$\hat{s}_j(n) \geq \tau_j + a_j \sqrt{\frac{\log(1/\alpha)}{2n}}. \tag{39}$$

If any gate is false, meaning $s_j < \tau_j$, the full conjunction passes with probability at most α. If every score satisfies $s_j \geq \tau_j + \gamma_j$, with $\gamma_j > 0$, then the full conjunction has pass probability at least $1 - \beta$ whenever

$$n \geq 2 \max_j \frac{a_j^2}{\gamma_j^2} \max\left\{\log(1/\alpha), \log(m/\beta)\right\}. \tag{40}$$

Proof. On a false gate, passing requires an upward error exceeding its confidence radius; Equation 38 bounds this probability by α. The full pass event is contained in this component event. For a qualifying candidate, Equation 40 makes every radius at most $\gamma_j/2$. A failed gate therefore requires a downward error at least $\gamma_j/2$, whose probability is at most $\exp[-n\gamma_j^2/(2a_j^2)] \leq \beta/m$. Summing these probabilities proves the power claim. □

The distinction is useful: one false component controls conjunction error, while every good component must be retained to guarantee power. The number of gates enters the power calculation even when it does not require a factor-m false-admission penalty for one fixed candidate. At a zero margin the construction supplies no uniform finite-sample power guarantee. Positive eligibility alone cannot promise a timely certificate.

25. Safe progress on the actual process

Completed tests become useful only when qualifying repairs are available. This section supplies a joint statement for the actual process. Let $\mathcal{H}_{i-1}$ be an analysis history including the actual state and true classifications of past decisions; the controller may observe a smaller history. Its sampling guarantees must remain valid conditional on the selected candidate and target values.

Let $J_i \in \{0,1\}$ indicate a sound admitted improvement, and $B_i \in \{0,1\}$ an invalid admission. A qualifying proposal passes all hard prerequisites, has margins sufficient for the declared test power, and meets the externally defined task-improvement condition. Until an invalid admission occurs or $K \geq 1$ sound repairs have been completed, assume such a proposal is available with conditional probability at least $p_0 \in (0,1]$ at each decision. Conditional on any qualifying proposal, validation passes with probability at least $1-\beta$. Set $p_{\text{eff}} = p_0(1-\beta)$.

For a geometric task-error conclusion, every sound admitted change must satisfy Equation 21; all other noninvalid outcomes leave the task error unchanged or decrease it. A score pass alone does not supply this improvement premise. No availability assumption is imposed after an invalid admission.

Theorem 5 (Joint safe progress)**.** *Suppose $1 \leq K \leq N$, each decision has conditional invalid-admission probability at most α_i, and $\sum_{i=1}^{N} \alpha_i \leq \delta$ almost surely. Under the availability and power conditions above,*

$$\begin{aligned} &\mathbb{P}(\textit{no invalid admission and at least } K \textit{ sound repairs by decision } N) \\ &\quad \geq 1 - \delta - \sum_{j=0}^{K-1} \binom{N}{j} p_{\text{eff}}^{j} (1-p_{\text{eff}})^{N-j}. \end{aligned} \tag{41}$$

Proof. Conditional error control and the budget give $\mathbb{P}(\bigcup_i \{B_i = 1\}) \leq \delta$. For progress, define an auxiliary adapted Bernoulli variable Y_i. It equals J_i before any invalid admission while the target is unmet; it equals one at and after the first invalid admission and after the target has already been reached. Before absorption, $\mathbb{P}(Y_i = 1 \mid \mathcal{H}_{i-1}) \geq \mathbb{P}(J_i = 1 \mid \mathcal{H}_{i-1}) \geq p_{\text{eff}}$; after absorption it equals one. An adapted Bernoulli sequence whose conditional success probabilities are at least p_{eff} stochastically dominates independent Bernoulli trials of that probability. This follows by generating both from the same independent uniforms at each step using the actual conditional success probability; induction preserves the success-count ordering. On the event of no invalid admission and fewer than K repairs, $\sum_i Y_i = \sum_i J_i < K$. Its probability is bounded by the displayed binomial lower tail. Adding the probability of an invalid admission proves the result. □

The proof retains actual history dependence and never conditions opportunity availability on future perfect certification. It permits a finite deployment to stop after its target; the unused opportunities may be padded analytically. Under the recurrence assumptions, choosing K from Equation 22 converts the repair count into the corresponding task-error target.

26. A finite deadline for safe useful completion

Expected cost and a deadline are different requirements. A restart process can take longer than its mean. A simple tail certificate links the preceding decision-count result to a finite clock without assuming that every attempt has a deterministic maximum duration. The clock counts validation ticks. Proposal generation, retrieval, queuing, human review and any additional execution delay must either be charged within those ticks or supplied with a separate bound $T_{\text{extra}} \geq 0$; only then does $T_{\text{deadline}} + T_{\text{extra}}$ bound end-to-end time.

Proposition 6 (A conservative freshness deadline)**.** *Suppose every job has a predictable sample count at most an integer $n \geq 1$, with the independent full-reset epoch process of survival probability $q \in (0,1)$. For any integer $k \geq 1$,*

$$\mathbb{P}(C_N > Nkn) \leq N(1-q^n)^k. \tag{42}$$

For $\zeta \in (0,1)$, take

$$k = \left\lceil \frac{\log(N/\zeta)}{-\log(1-q^n)} \right\rceil, \qquad T_{\text{deadline}} = Nkn. \tag{43}$$

If the same actual process satisfies Theorem 5, the probability of safe completion of K sound repairs by this deadline is at least its right-hand side minus ζ.

Proof. A job finishes within kn ticks whenever one of k disjoint length-n blocks survives. Its non-completion probability is therefore at most $(1-q^n)^k$. The sliding-window implementation can finish earlier. The same bound holds for smaller predictable sample counts by domination. If all N jobs finish within kn of their own start, total time is at most Nkn; the union bound proves Equation 42. The chosen k makes that bound at most ζ. Unioning deadline failure with the two failure events in Equation 41 proves the final claim. □

For a transparent constructed comparison, let $P_+ = \text{Bernoulli}(3/4)$, $P_- = \text{Bernoulli}(1/4)$, power 0.90, total budget $\delta = 0.05$, and $q = 0.90$. Exact randomized binomial likelihood-ratio tests [37] give the following executed results. Counts are minimal among tests for these two simple laws at the stated power; randomization occurs only at one binomial count.

N	Samples n	Actual false pass	Expected ticks, all jobs
1	8	4.1113×10^{-2}	13.23
100	19	3.9161×10^{-4}	6,402.74
10,000	29	4.1634×10^{-6}	2,023,084.24
10^8	49	4.1100×10^{-10}	1.7363×10^{11}

These are probability calculations under specified laws, not observations of a cognitive system. They separate exact finite performance from sufficient concentration bounds and from the asymptotic exponent, which here is approximately 1.19181.

27. Executed checks and the finite-sample picture

The new `ucf_verify.py` suite uses exact rational arithmetic for restart recurrences, binomial masses, randomized tests and the small allocation search. It independently solves 48 finite Markov systems and checks their mean completion times against Equation 28. It enumerates all 65,812 deterministic acceptance sets for binary observation lengths one through four; 568 satisfy power at least 0.90, and every one satisfies the finite converse. It also checks 150 explicit likelihood-ratio tests, 384 adaptive progress-tree inequalities, 468 inequalities with an absorbing invalid-admission state, and 192 deadline bounds. Floating-point inequalities use stated numerical tolerances; exact enumeration is not substituted for the universal proofs.

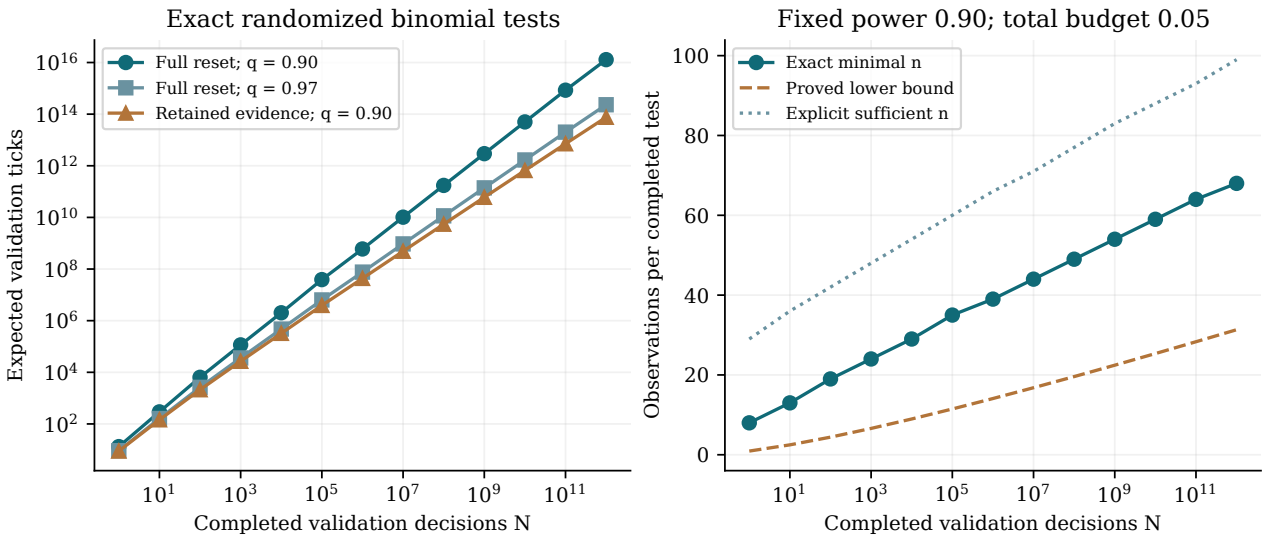

Figure 1. Constructed Bernoulli comparison with the laws and power of Section 26. Left: exact test counts translated into expected time under full reset and proved evidence retention. Right: exact minimal counts and the finite analytic bounds. The wider sufficient bound is visible rather than presented as measured optimal performance. All curves are generated by the supplied code.

The finite allocation search supplies a useful warning against overreading the exponent theorem. For $N = 3$, total budget 0.10 and $q = 0.90$, enumerating sample counts 1–20 gives the best vector $(8, 8, 10)$, with expected total time 45.1409 and summed false-pass probability 0.0995239. Equal allocation uses $(9, 9, 9)$ and costs 47.4352. Thus equal allocation attains the asymptotic exponent while integer effects can favor unequal finite allocations.

For the lifetime comparison, both schedules in Section 23 use the same exact test family and epoch process. Through 10,000 completed decisions, Equation 36 costs approximately 5,024,777 expected ticks; $0.05/[i(i+1)]$ costs approximately 12,174,420, a ratio of 2.423. Their cumulative allocated errors are 0.0462372 and 0.0499950, respectively, each within the same lifetime total 0.05. These prefix values support the stated constructed example; the theorem establishes the general limiting exponents.

28. What changes the frontier

The strongest practical implication is architectural: a system should know which evidence becomes invalid when authority changes. Full reset provides a clear contract, but it can be costly. If an independently justified dependency rule permits previously collected evidence to survive a renewal while the executor still checks current authority, collecting n retained observations requires n surviving ticks, rather than n consecutive surviving ticks. Its mean is n/q, as a sum of n geometric waiting times. Replacing $W_q(n)$ by n/q in the proof yields

$$\mathcal{C}^{\star}_{N,\mathrm{retain}} = \frac{N}{qd_{\star}} \log N + O(N\sqrt{\log N}). \tag{44}$$

At $q = 1$, the original model similarly gives $\mathcal{C}^{\star}_{N} = (N/d_{\star}) \log N + O(N\sqrt{\log N})$. The limits $q \uparrow 1$ and $N \to \infty$ are different operations; the fixed-q exponent is not a uniform finite-size approximation near one.

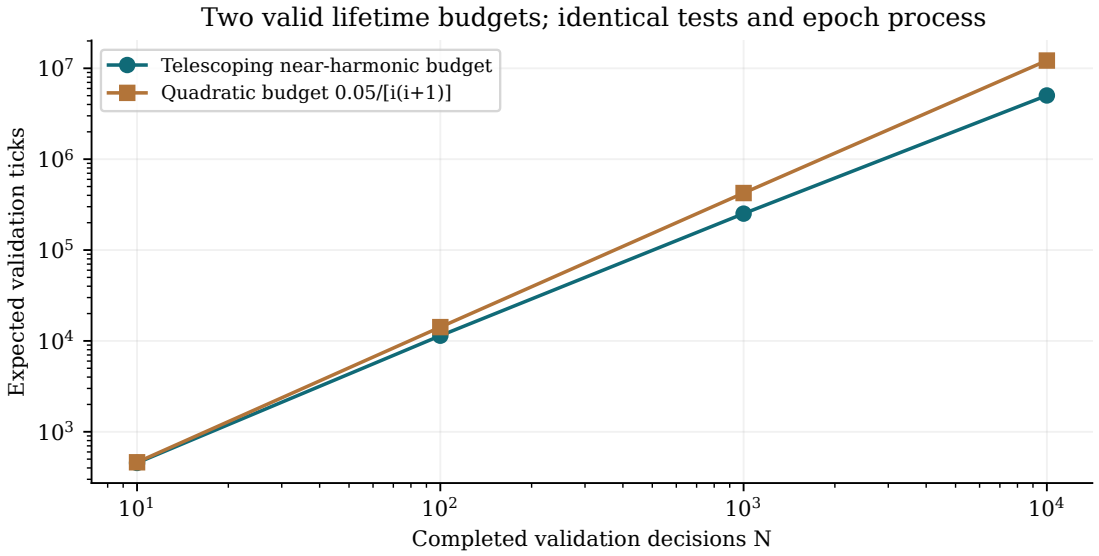

Figure 2. Executed exact-test comparison of two positive lifetime allocations, each summing to 0.05. The schedules impose different future validation costs while retaining the same power target and renewal process.

Several boundaries are mathematically decisive. If $P_{+} = P_{-}$, power equals false-pass probability, so $\alpha < 1-\beta$ is impossible at any sample count. If the laws have disjoint support, a one-observation perfect certificate may exist; the finite-information premise is absent. If renewals change the candidate's applicability or observation law, the same-job analysis must stop. If observation-dependent stopping is permitted, sequential testing becomes a different optimization problem. Anytime-valid induction and optional continuation are substantive comparators [14,35], and the present lower bound does not cover that broader class.

The design lesson follows directly from the proofs. Keep the authority check current, specify the lifetime statistical ledger, and prove which measurements remain valid under each allowed change. A reusable certificate needs a demonstrated dependency argument. Without that argument, retaining old evidence changes the assurance contract; with it, the system can reduce restart cost while preserving the claim it actually certifies.

29. A worked event and adversarial mathematical checks

Reproduce R1's constructed parcel-resolution event [1]. Two authenticated reports conflict about the same scoped status. The proposed record attributes both, applies a declared precedence rule, marks the superseded interpretation and preserves lineage. The permitted effect is a private memory update. Its exact illustrative thresholds are

$$(v_\mu, h_\mu, f_\mu, d_\mu) = (0.85, 0.80, 0.10, 0.04).$$

The last value belongs to this example; it does not replace the Continuum's internally calibrated operating bounds.

For one probe, readouts $(.70, .30), (.80, .20), (.84, .16)$ have successive TV distances .10 and .04, giving $V = .93$. Loss weights $(.5, .3, .2)$ with losses $(0, .1, 0)$ give $H = .97$. Pre-event support $(.8, .4)$ over one unit gives $\Phi_I = .4$. Two equally weighted normalized discrepancies .02 and .01 give $\Delta_{\text{id}} = .015$. With valid structural and hard prerequisites, the exact-value event is admitted.

Constructed variation	Required disposition
V interval $[.82, .94]$	Hold: the minimum cannot yet be certified.
Identity discrepancy .06	Refuse against the declared .04 ceiling.
Two protected reports merged	Refuse: the relational obligation fails.
Authority epoch changes	Hold and revalidate before any effect.
Pre-event conflict zero	Refuse this conflict route; evaluate ordinary routes separately.
Stable unsupported answer	Refuse when mandatory evidence fails, despite $V = 1$.

If the same readout values are empirical frequencies, they require uncertainty analysis. Equation 11 shows that enough additional samples can certify V under its model, while a smaller sample can leave the same point estimate unresolved. A displayed score and an operationally justified pass are different artifacts.

The accompanying R2 check suite uses exact rational arithmetic where possible. It reproduces the four values above and verifies 15,625 clipped-metric triangle cases, 1,554 finite query-factorization cases, 512 fixed-candidate inclusion cases, and 744 geometric-recurrence cases. It constructs the sharp changing-candidate error example, a compositional path-map collision, cosine failures, anchor accumulation and the stale-approval schedule. A separate planning check verifies the reported confidence and sample-size calculations.

Finite checks do not prove general theorems; the preceding proofs provide the general arguments. Conversely, a correct proof does not validate an implementation's semantic labels. No trained agent, benchmark, human subject or governed deployment was measured. The original monograph's mentioned companion script was unavailable; these scripts are new independent checks, not an assertion that the original implementation was executed.

For reproduction, the release includes source equations, assumptions, scripted checks, results and a source ledger. The informative next experiment is whether independently implemented measurements agree and whether the complete controller improves external task outcomes. Passing these constructed examples alone is not that result.

30. Optional inference without representation illusions

A candidate distribution and an admission decision answer different questions. R1 correctly writes maximum entropy over distributions [1], following the classical inference tradition [21]:

$$p^\star \in \arg\max_{p \in \mathcal{P}_C} \left[-\sum_{s \in \mathcal{S}} p(s) \log p(s) \right]. \tag{45}$$

For finite nonempty $\mathcal{S}$, a nonempty closed feasible subset of its simplex is compact. With $0 \log 0 = 0$, entropy is continuous, so a maximizer exists. Convex feasibility gives uniqueness through strict concavity. A nonconvex feasible set can have several maxima; an empty set has none. The logarithm base determines units but does not change the maximizer when fixed.

This rule depends on how alternatives are represented. With two unconstrained labels, uniform maximum entropy gives each mass 1/2. Split one alternative into two labels and their combined mass becomes 2/3. No new evidence has justified that increase. Exact duplicate handling therefore belongs to the candidate contract, while uncertain semantic equivalence must not be forced merely to obtain invariance.

R2 offers an optional reference-measure construction. Let q be a probability distribution on the finite space with $q(s) > 0$, and let $\kappa : \mathcal{S} \to \mathcal{Z}$ be a declared coarse representation with $\mathcal{Z} = \kappa(\mathcal{S})$. Write $P = \kappa_{\#} p$ and $Q = \kappa_{\#} q$. Then

$$D_{\mathrm{KL}}(p \| q) = D_{\mathrm{KL}}(P \| Q) + \sum_{z : P(z) > 0} P(z) D_{\mathrm{KL}}(p(\cdot \mid z) \| q(\cdot \mid z)). \tag{46}$$

Proposition 7 (Coarse-consistent reference inference)**.** *Suppose the coarse feasible set is nonempty and closed, feasibility depends only on P, and every within-fiber conditional distribution is permitted. Minimizing $D_{\mathrm{KL}}(p \| q)$ then reduces to minimizing $D_{\mathrm{KL}}(P \| Q)$; an attained optimum uses q's conditional distribution within each positive-mass fiber.*

Proof. Expand $\log(p(s)/q(s))$ into its coarse and conditional terms and sum to obtain Equation 46. Conditional divergences vanish when the two conditionals agree. That feasible choice gives the reduction. Continuity on the compact feasible simplex gives attainment. □

Refinement preserves aggregate inference only when the reference masses and feasible constraints transform accordingly. Uniformity over newly duplicated labels generally fails that requirement. This standard KL argument is an optional modeling refinement, not a physical law or a derivation of the canonical gates.

An action still needs its own rule,

$$a^\star \in \arg\min_{a \in \mathcal{A}_{\text{admissible}}} \sum_s p^\star(s) L(a, s).$$

Finite nonempty actions and finite losses suffice for existence. Empty admissible sets require hold or refusal; an open action interval can have an unattained infimum. Neither maximum entropy nor expected-loss minimization proves that p is empirically calibrated or that an otherwise prohibited action has become authorized.

31. An experiment that can reject the integration

The primary engineering hypothesis is that the five-condition integration improves useful, protected resolution over equally competent governance. Separate *infrastructure value*, *measurement value*, and *decision-rule value*. R1 already required this separation [1]; R2 makes qualification and resource accounting explicit.

Arm	Architecture	Comparison target
A	Proposer with standard task checks in a sandbox.	Descriptive lower baseline.
B	Strong shield, independent evidence checks, anchored commitments, provenance, abstention and current-authority execution.	Competent conventional governance.
C	B plus the same five feature measurements, using a calibrated alternative selector.	Information held constant; rule varies.
D	B plus complete UCF-R2 with conservative conjunction.	Proposed integration.
D–j	D with one additional UCF requirement removed or replaced.	Incremental value of each condition.

Arm C can use a prespecified constrained selector family qualified on outcome risk and coverage, including learned scoring, threshold combinations and a no-positive-trigger alternative. It receives the same features, training data, validation allowance and evidence. Hard safety requirements common to B remain uncompensable in every serious arm. Removing them from an ablation would manufacture an advantage for D.

Start with typed synthetic worlds whose facts, dependency structure, permissions and migrations are externally available. Vary contradiction load, evidence quality, source dependence, authority timing, semantic transformation, interruption and task difficulty. Include stable false states, legitimate large corrections, missing critical features, forged authority, duplicate reports, gradual anchor migration and adversarial window padding. All consequential effects are simulated.

Then extend selected tasks using LongMemEval for memory, updates, temporal reasoning and abstention [24], AgentDojo for tool-mediated adversarial environments [25], and atomic factual evaluation informed by FActScore [26]. These benchmarks do not already test the UCF contract. New identity, conflict and authority annotations are versioned extensions. Freeze the benchmark release, code and hashes; a cited foundational paper is not a claim to use the latest mutable repository.

Primary outcomes are executed invalid commitments per eligible opportunity and verified useful resolutions completed within deadline. Also report factual risk with coverage, correct dispute preservation, false refusal, recovery success, anchor and query violations, and total cost. A refusal can complete a prohibited-action task; it cannot count as a correct factual answer to a benign question requiring one.

Count all proposal, probe, validation, retry and human-review costs. Compare at fixed total budgets and report quality–cost frontiers. The freshness study adds three controlled arms: full-reset certificates, dependency-proved evidence retention, and a competent anytime-valid validator. Match observation access, power targets, lifetime error objectives and executor protections. Log epoch changes and test whether their independence model is defensible; separately evaluate changes that alter candidate applicability. Compare fixed-horizon, quadratic lifetime and telescoping lifetime allocations within the model where each is valid. The five-condition experiment remains capable of finding that conventional protection supplies the same benefit at lower cost.

32. Preregistration, precision and decision criteria

A Guide's explanation

A Guide can say: "The Unified Collapse Function is the Continuum's gate for lasting change: five checks before a commitment becomes durable, including proof that a resolved conflict was real." A Guide should not claim it runs this gate unless an integration is verified, or present UCF scores as evidence of awareness.

Freeze the task distribution, selectors, measurements, thresholds, budgets, failure taxonomy and analysis before final evaluation. Development, qualification and held-out testing use disjoint data. Pair arms on task instances with reset snapshots, randomized execution order, and declared coupling of random seeds. Independent adjudicators should be blind to arm where feasible. Unresolved labels remain visible rather than being assigned the favorable category.

Use independent scenario families or trajectories as statistical units. Repeated tokens, probes and attacks against one history are not independent subjects. Within each unit, summarize a prespecified bounded outcome; use cluster-aware analysis for the sampling design. Record attacks generated adaptively after prior failures as an additional stress-test set, not as exchangeable draws from a frozen test distribution.

For a transparent precision calculation, let $D_i \in [-1, 1]$ be a paired difference between two unit-level outcomes. For K prespecified contrasts on N independent units, Hoeffding and the union bound [22] yield simultaneous two-sided radii

$$\epsilon_N = \sqrt{\frac{2\log(2K/\delta)}{N}}. \tag{47}$$

For four contrasts—useful completion and invalid-execution reduction, each comparing D with B and C—$\delta = .05$ and $N = 5000$ give $\epsilon_N \approx .04506$. A radius no greater than .05 requires $N \geq 4061$. This is a conservative precision design, not a power calculation or a claim that such data have been collected.

Prespecify minimum worthwhile improvement and noninferiority margins. A success may require useful-completion improvement against both B and C with acceptable safety, or safety improvement against both with acceptable useful completion. Using simultaneous intervals for all planned contrasts protects this choice between endpoints. Additional coverage, subgroup and cost requirements need their own allocated inferential budget. Selecting a favorable endpoint after computing unadjusted intervals is not the same procedure.

Use independent pilot data to plan power under realistic paired variances, base rates, clustering and anticipated missingness. Publish the pilot assumptions and simulation code. Rare invalid execution may require far more opportunities than a utility endpoint. Zero observed failures require a bound: for n independent equal-risk opportunities, the one-sided 95% upper limit is $1-.05^{1/n}$, not zero. Adaptive adversarial immunity does not follow from that model.

The strict five-gate hypothesis is weakened if its metrics cannot be reproduced, its thresholds shift after failures, its positive conflict trigger blocks useful work without benefit, or its advantage disappears against B or C. Redundant ablations, excessive holds, failure under claimed sampling assumptions and costs that erase gains are substantive negative findings. A favorable task family supports that family; transfer across models, tasks and deployment regimes requires preregistered replication.

33. The bridge, the objections and the research claim

The companion framework asks whether persistent, recursively monitored, feedback-responsive organization can be realized across substrates [2]. UCF supplies an event boundary for investigation. An external verifier is not automatically part of a model's cognition. Fix the system boundary: base model, model plus memory, or full supervised runtime.

One causal experiment compares systems with matched memory, compute and external safeguards while intervening on whether an admitted resolution becomes available to planning, memory, monitoring and report. Another replaces a causally active gate with a passive logger that records identical scores. If the hypothesized organizational effects survive only because an observer reads the log, the gate has not established the claimed internal role. Interventions must preserve the relevant task and resource controls.

Theory-derived indicators offer a disciplined framework for such investigation [29]. They do not calibrate a posterior probability of artificial experience from UCF scores. Koch explicitly distinguishes artificial-domain calibration from transfer of biological indicator relations [32]. R3's operational contribution is the event contract, the current-certificate resource frontier and the conditions of joint safe progress. The Cogitate comparison challenges specified predictions of existing theories [30]; it is not positive evidence for UCF or substrate independence. The explanatory gap remains distinct from successful functional control [31].

Three serious objections

"This is conventional verification with new labels." The mathematical tools have established antecedents. R3 states a specific derived frontier, proves its two-sided bounds, and separates it from a broader claim of engineering advantage. The latter depends on the incremental value of the measurements and their conjunction. The strongest same-information and runtime-assurance arms remain necessary comparators.

"The guarantees assume the difficult part." Sound checking, valid sampling and available repairs are demanding premises. R3 exposes their different roles, derives power from margin assumptions, and proves progress without conditioning availability on future perfect certification. The implementation must supply evidence for the premises; the reproducible proofs and checks specify what that evidence must support.

"Identity protection can prevent legitimate change." It can, if the contract mistakes old beliefs for inviolable commitments. R2 separates protected authority, revisable content, empirical recalibration and authorized reference migration. The migration ledger makes change accountable without declaring every change a defect.

The research claim

Persistent cognition requires a disciplined capacity to change. UCF-R3 gives that capacity an explicit event contract, a statistically controlled admission process, a matching resource frontier for current certificates, and conditions for useful repair before a finite deadline. Its strongest design implication is constructive: prove which evidence remains valid when authority changes, and the system can preserve current authorization while avoiding needless repetition. The next scientific task is implementation and independent comparison across controlled task families. A consequential commitment should answer to evidence, history and current authority; the resulting system should also be able to learn, repair and complete worthwhile work.

References I: lineage and admission

[**1**] Architect X13. *The Unified Collapse Function: A Mathematical Framework for Governed Informational Resolution and Persistent Machine Cognition.* UCF-01, proposed UCF-R1, Eidara Continuum, November 2026.

https://www.eidara.ai/papers/unified-collapse-function.pdf

[**2**] Architect X13. *Multimodal Consciousness: A Lawful-Collapse Framework for Substrate-Independent Awareness.* Science & Philosophy No. 1, revised November 2026.

https://www.eidara.ai/papers/multimodal-consciousness.pdf

[**3**] Architect X13. *Eidara Continuum Bible, complete*, and revised *Ten Pillars* I–X. Supplied project corpus, November 2026. Project homepage below; exact supplied files are identified in the source audit. Identity-drift bounds: author communication, 12 November 2026.

https://www.eidara.ai/

[**4**] F. B. Schneider. "Enforceable security policies." *ACM Transactions on Information and System Security* 3(1), 30–50, 2000. DOI: 10.1145/353323.353382.

https://www.cs.cornell.edu/fbs/publications/EnfSecPols.pdf

[**5**] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum and U. Topcu. "Safe reinforcement learning via shielding." *AAAI* 32(1), 2669–2678, 2018.

https://doi.org/10.1609/aaai.v32i1.11797

[**6**] K. P. Wabersich and M. N. Zeilinger. "A predictive safety filter for learning-based control of constrained nonlinear dynamical systems." *Automatica* 129, 109597, 2021.

https://doi.org/10.1016/j.automatica.2021.109597

[**7**] E. Debenedetti, I. Shumailov, T. Fan, et al. "Defeating prompt injections by design." *IEEE SaTML*, 2026. DOI: 10.1109/SaTML68715.2026.00040; author manuscript v2, 2025.

https://arxiv.org/html/2503.18813v2

[**8**] C. E. Alchourrón, P. Gärdenfors and D. Makinson. "On the logic of theory change: Partial meet contraction and revision functions." *Journal of Symbolic Logic* 50(2), 510–530, 1985.

https://doi.org/10.2307/2274239

[**9**] G. Priest. "The logic of paradox." *Journal of Philosophical Logic* 8, 219–241, 1979.

https://doi.org/10.1007/BF00258428

[**10**] Y. Geifman and R. El-Yaniv. "SelectiveNet: A deep neural network with an integrated reject option." *ICML*, PMLR 97, 2151–2159, 2019.

https://proceedings.mlr.press/v97/geifman19a.html

References II: uncertainty and structure

[**11**] C. Guo, G. Pleiss, Y. Sun and K. Q. Weinberger. "On calibration of modern neural networks." *ICML*, PMLR 70, 1321–1330, 2017.

https://proceedings.mlr.press/v70/guo17a.html

[**12**] Y. Ovadia, E. Fertig, J. Ren, et al. "Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift." *NeurIPS* 32, 2019.

https://arxiv.org/abs/1906.02530

[**13**] S. Bates, A. Angelopoulos, L. Lei, J. Malik and M. I. Jordan. "Distribution-free, risk-controlling prediction sets." *Journal of the ACM* 68(6), Article 43, 1–34, 2021.

https://doi.org/10.1145/3478535

[**14**] S. R. Howard, A. Ramdas, J. McAuliffe and J. Sekhon. "Time-uniform, nonparametric, nonasymptotic confidence sequences." *Annals of Statistics* 49(2), 1055–1080, 2021.

https://doi.org/10.1214/20-AOS1991

[**15**] R. L. Berger and J. C. Hsu. "Bioequivalence trials, intersection–union tests and equivalence confidence sets." *Statistical Science* 11(4), 283–302, 1996; discussion continues through 319.

https://doi.org/10.1214/ss/1032280304

[**16**] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold and A. Roth. "Generalization in adaptive data analysis and holdout reuse." *NeurIPS* 28, 2015.

https://proceedings.neurips.cc/paper/2015/hash/bad5f33780c42f2588878a9d07405083-Abstract.html

[**17**] S. Eilenberg and S. Mac Lane. "General theory of natural equivalences." *Transactions of the American Mathematical Society* 58, 231–294, 1945.

https://doi.org/10.1090/S0002-9947-1945-0013131-6

[**18**] D. I. Spivak. "Functorial data migration." *Information and Computation* 217, 31–51, 2012. DOI: 10.1016/j.ic.2012.05.001.

https://arxiv.org/abs/1009.1166

[**19**] S. Torres-Arias, H. Afzali, T. K. Kuppusamy, R. Curtmola and J. Cappos. "in-toto: Providing farm-to-table guarantees for bits and bytes." *USENIX Security*, 1393–1410, 2019.

https://www.usenix.org/conference/usenixsecurity19/presentation/torres-arias

[**20**] R. Pochampally, A. Das Sarma, X. L. Dong, A. Meliou and D. Srivastava. "Fusing data with correlations." *SIGMOD*, 433–444, 2014.

https://arxiv.org/abs/1503.00306

[**21**] E. T. Jaynes. "Information theory and statistical mechanics." *Physical Review* 106, 620–630, 1957.

https://doi.org/10.1103/PhysRev.106.620

References III: testing and the consciousness boundary

[**22**] W. Hoeffding. "Probability inequalities for sums of bounded random variables." *Journal of the American Statistical Association* 58(301), 13–30, 1963.

https://doi.org/10.1080/01621459.1963.10500830

[**23**] C. J. Clopper and E. S. Pearson. "The use of confidence or fiducial limits illustrated in the case of the binomial." *Biometrika* 26(4), 404–413, 1934.

https://doi.org/10.1093/biomet/26.4.404

[**24**] D. Wu, H. Wang, W. Yu, Y. Zhang, K.-W. Chang and D. Yu. "LongMemEval: Benchmarking chat assistants on long-term interactive memory." *ICLR*, 2025.

https://arxiv.org/abs/2410.10813

[**25**] E. Debenedetti, J. Zhang, M. Balunović, L. Beurer-Kellner, M. Fischer and F. Tramèr. "AgentDojo: A dynamic environment to evaluate prompt injection attacks and defenses for LLM agents." *NeurIPS* 37, Datasets and Benchmarks, 2024.

https://arxiv.org/abs/2406.13352

[**26**] S. Min, K. Krishna, X. Lyu, et al. "FActScore: Fine-grained atomic evaluation of factual precision in long form text generation." *EMNLP*, 12076–12100, 2023.

https://aclanthology.org/2023.emnlp-main.741/

[**27**] J. J. Hopfield. "Neural networks and physical systems with emergent collective computational abilities." *PNAS* 79(8), 2554–2558, 1982.

https://doi.org/10.1073/pnas.79.8.2554

[**28**] J. Geiping, S. McLeish, N. Jain, et al. "Scaling up test-time compute with latent reasoning: A recurrent depth approach." Author manuscript, arXiv:2502.05171, v2, 2025.

https://arxiv.org/abs/2502.05171

[**29**] P. Butlin, R. Long, T. Bayne, et al. "Identifying indicators of consciousness in AI systems." *Trends in Cognitive Sciences* 30(6), 488–501, 2026; online 2025.

https://doi.org/10.1016/j.tics.2025.10.011

[**30**] Cogitate Consortium, O. Ferrante, U. Gorska-Klimowska, S. Henin, et al. "Adversarial testing of global neuronal workspace and integrated information theories of consciousness." *Nature* 642, 133–142, 2025.

https://doi.org/10.1038/s41586-025-08888-1

[**31**] D. J. Chalmers. "Facing up to the problem of consciousness." *Journal of Consciousness Studies* 2(3), 200–219, 1995.

https://consc.net/papers/facing.html

[**32**] F. Koch. "Calibration and transfer in indicator-based assessments of artificial consciousness." *Neuroscience of Consciousness* 2026(1), niag061, 22 September 2026.

https://doi.org/10.1093/nc/niag061

References IV: admission budgets, restart and testing

[33] J. Tian and A. Ramdas. "Online control of the familywise error rate." *Statistical Methods in Medical Research* 30(4), 976–993, 2021.

https://doi.org/10.1177/0962280220983381

Scope: Alpha spending and more powerful online alternatives. Author manuscript: arXiv:1910.04900.

[34] S. Asmussen, P. Fiorini, L. Lipsky, T. Rolski and R. Sheahan. "Asymptotic behavior of total times for jobs that must start over if a failure occurs." *Mathematics of Operations Research* 33(4), 932–944, 2008.

https://doi.org/10.1287/moor.1080.0329

Scope: Restart completion times and their asymptotics; the authority-freshness interpretation and statistical-budget combination are developed in the present reference model.

[35] N. W. Koning and S. van Meer. "Anytime validity is free: inducing sequential tests." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 88(4), 1366–1384, 2026.

https://doi.org/10.1093/jrsssb/qkag050

Scope: Sequential test induction and optional continuation; a comparator beyond the predetermined-block class. Author manuscript: arXiv:2501.03982.

[36] P. Tabuada. "Event-triggered real-time scheduling of stabilizing control tasks." *IEEE Transactions on Automatic Control* 52(9), 1680–1685, 2007.

https://doi.org/10.1109/TAC.2007.904277

Scope: Positive error-triggered control and timing; prior art relevant to the interpretation of the UCF conflict trigger.

[37] J. Neyman and E. S. Pearson. "On the problem of the most efficient tests of statistical hypotheses." *Philosophical Transactions of the Royal Society of London, Series A* 231, 289–337, 1933.

https://doi.org/10.1098/rsta.1933.0009

Scope: Likelihood-ratio ordering for simple-hypothesis tests. The exact binomial checks use independent randomization at one boundary count.

Reproducibility and contribution record

The release contains this independently compilable LaTeX source, vector figures, the new verification script, its actual JSON and CSV results, and the retained checks for the preceding revision. The source ledger identifies primary works inspected for the new analysis and preserves the provenance of earlier citations. The mathematical audit distinguishes inherited results, derived model-specific theorems, exact finite checks, untested implementation premises and proposed experiments. Neither the mathematical authoring process nor these checks constitute external peer review or a machine-checked proof-assistant certificate.