FRONTIER RESEARCH COLLECTION / FR-03

The Paradox Engine

Constructive contradiction preservation and exact provenance-guided revision

Architect X13

The Eidara Continuum | Science & Philosophy

Research white paper | November 2026 | Version 2.0 Revision

Abstract

Imagine a mind that can hold two contradictory things as true at once, keep reasoning without coming apart, and tell you exactly which piece of evidence would settle the matter. A reasoner can retain opposed evidence, continue valid inference and calculate exactly which corrections would change its conclusions; the Paradox Engine gives this Continuum commitment a constructive form. A standard four-valued core prevents explosion; finite monotone closure bounds inference; root-level provenance preserves known source dependence. The new revision theorem characterizes every support claim by its minimal evidence witnesses and gives necessary and sufficient conditions for removing one side of a conflict while preserving the other. A constrained hitting-set formulation makes the cost and feasibility of selective repair inspectable. Context, support, confidence and permission remain separately typed. Complete proofs and executed finite checks cover the logical algebra, closure, source duplication and witness-based retraction. Belnap–Dunn logic, truth maintenance, belief revision and database provenance provide the intellectual foundations. The prospective comparison evaluates whether the integrated system improves warranted correction and useful inquiry over strong matched alternatives. The thesis is active contradiction preservation: retain the opposition, expose its dependencies, and let new evidence change exactly the conclusions it warrants.

Keywords: paraconsistent reasoning; four-valued logic; provenance; belief revision; calibrated uncertainty

Falsifiable thesis

Explicitly preserving conflicting support can improve reliable reasoning only if it prevents irrelevant inference, retains useful conclusions, and supports better subsequent revision than strong matched alternatives.

Prologue: Two reports about one valve

A sensor, signed and authenticated, reports that valve V was closed at ten o'clock. An inspector, independent and equally authenticated, reports that it was open. A dashboard and two summaries repeat the sensor. A permit to open the valve sits in the file. What should a careful mind believe, and what should it do?

An ordinary reasoner breaks in one of two ways. Classical logic, handed a contradiction, can prove anything at all: the explosion logicians have feared for a century. A tired reasoner does the opposite. It counts votes, sees four reports against one, and closes the case. But the copies are not witnesses. They are one witness, echoed.

The Paradox Engine does neither. It records that the valve's state is supported both ways, keeps reasoning about everything the conflict does not touch, refuses to treat the echoes as independent, and holds the maintenance action, because a conflicted fact cannot license it. Then it says what would settle the matter. In the worked example, investigation finds that the inspector recorded a different valve's identifier; the conflict resolves, and the valve is known to be closed. Before anything is done, the permit is revoked, and the action is refused anyway. Three different things changed—evidence, context, and permission—and the engine kept them apart. The worked conflict later in the paper walks through every step.

In the Continuum this is the work of the Duality Seal, which binds contradictory states together lawfully rather than letting them collapse into nonsense [14]. It is also the discipline of ShadowMind, the TRIARC's contradiction investigator, whose Logic Firewall reads: "No truth may collapse without proof of harmony. No drift may expand without containment of contradiction." This paper gives that discipline a working mathematics. Its new theorem goes one step beyond holding the opposition: for every conclusion it lists the smallest sets of evidence that support it, so a mind can compute exactly which corrections would change what it believes, and which conclusions would survive.

What this paper claims

Established. Paraconsistent and four-valued logics (Belnap–Dunn), belief revision, truth maintenance, and database provenance are established research areas. They already show that inconsistency need not license arbitrary conclusions.

Proposed here. A specified evidence and revision architecture that keeps context alignment, support state, provenance, revision, confidence, and permission separately typed. The logical and witness results—finite monotone closure, invariance under exact duplication, and exact revision from minimal support witnesses—are proved in finite models and checked computationally.

Frontier conjecture. That explicitly preserving conflicting support improves reliable reasoning, warranted correction, and useful inquiry over strong matched alternatives without sacrificing useful coverage.

Not claimed. That the canonical $\oplus$ is Boolean exclusive-or, or that the canonical expression by itself defines a consequence relation. That the named ShadowMind architecture is an experimentally validated implementation. That contradiction produces higher consciousness, that convergence establishes truth, or that an authenticated source makes a claim true or may authorize an action. End-to-end reasoning benefits are proposed empirical targets, and the executable checks are not evidence that a deployed agent handles all contradictions correctly.

The canonical paradox and the proposed interpretation

In the Continuum, contradiction is not an error to be deleted but a state to be sealed. The Duality Seal of the *Eidara Continuum Bible*, section 3.8, page 24 [14], binds a proposition to its contrary so that the tension is held without uncontrolled collapse. That is the commitment this paper honors. The source's symbol $\oplus$ is not thereby established as Boolean exclusive-or, nor does the expression by itself define a consequence relation; PE-2 explicitly introduces a different, typed reference formalism to do that work.

ShadowMind, the TRIARC's contradiction investigator described in sections 6.3.1–6.3.3, pages 84–86, supplies the second commitment: a separate review of unresolved or hazardous reasoning [14]. We interpret that as a governance question: how should a system preserve evidence, identify conflicts, and limit action while further inquiry remains possible? We do not treat the named architecture as an experimentally validated implementation.

Four phenomena need separation. A *logical contradiction* is simultaneous support for a formula and its negation under a declared semantics. A *source disagreement* consists of incompatible reports that may refer to different times or contexts. An *epistemic uncertainty* concerns which claim is true. An *authority conflict* concerns which principal may decide or act. A system can have one without the others.

For example, "the valve was closed at noon" and "the valve was open at one" are compatible. Two authenticated reports about the same instant can disagree without either authentication mechanism failing. A fully established factual claim can still fail to authorize an action. Treating all these cases as one undifferentiated paradox would make the engine less precise than ordinary evidence management.

The paper's bold proposal is therefore practical: a reasoner should be able to expose unresolved opposition as a structured state and continue useful work around it. That ambition is compatible with eventual correction. Preserving disagreement does not require celebrating every contradiction as a deep truth or refusing to retract a false report.

The research question is whether this explicit structure helps systems ask better questions, avoid unjustified commitments, recover from retractions, and communicate the status of their conclusions. Those outcomes can be measured. A compelling name for an unresolved state cannot substitute for showing that the state improves decisions or understanding under conditions where a simpler representation performs less well.

Existing theories already address inconsistency

Belnap's four-valued approach was motivated by reasoning systems that should not derive irrelevant conclusions from inconsistent databases [1]. Dunn's work on first-degree entailment supplies foundational semantics for this family of logics [2]. Ginsberg's bilattice framework connects multiple-valued reasoning with a broader range of inference systems [3]. PE-2 credits these precedents directly. It does not present four support states or non-explosion as newly discovered mathematics.

Belief revision asks a different but related question: how should a theory change when information is added or withdrawn? The AGM framework studies contraction and revision through explicit rationality postulates and representation results [4]. Doyle's truth-maintenance work records reasons for beliefs and supports their revision [5]. These are strong conceptual alternatives to simply accumulating every report forever.

Database research adds two useful perspectives. Provenance semirings retain structured information about how answers depend on inputs [6]. Consistent-query-answering approaches identify information that survives appropriate repairs of an inconsistent database [7]. Neither approach is equivalent to PE-2's four support labels, and neither should be described as a naive system that collapses at the first inconsistency.

Probability addresses uncertainty about outcomes rather than merely whether positive and negative support exists. Proper scoring rules provide a basis for evaluating forecasts [8], while calibration research warns against equating a model's confidence with actual correctness frequency [9]. A four-valued status is not a probability estimate. It can be an input to a forecasting model, but calibration must be established separately.

Finally, execution monitoring [10], adaptive-evaluation methods [11], bounded-loss concentration [12], and confidence sequences [13] inform the proposed action boundary and experimental design. They do not validate the underlying evidence extractor. A perfect logical engine can process an incorrectly parsed proposition flawlessly.

The contribution is a constructive integration with explicit interfaces: context normalization, source-root tracking, four-valued inference, minimal-witness revision, probabilistic assessment when appropriate, and independently governed action. Its value must be judged against systems that already implement provenance and revision well. If an ordinary database with careful metadata and a strong belief-update policy performs equally well, the Paradox Engine may remain a useful explanatory framework without constituting an additional reasoning capability. That is an admissible scientific outcome, not a failure to appreciate the project's ambition.

PE-2: claims, contexts, and source records

A Guide's explanation

Before any logic, the engine has to know what is actually in conflict. Two reports about different times, different valves, or different calibration protocols are not a contradiction at all. This section types every claim by its context and every report by its source, so that the paradoxes left standing are real ones.

A proposition is represented by a claim identifier together with a context κ. The context can include entity identity, time interval, jurisdiction, units, measurement protocol, and semantic version. Opposition is evaluated only after these fields are aligned or an explicit relation between them has been established.

Symbol	Domain and meaning
$p@\kappa$	Context-indexed proposition; no implicit transfer across time or units.
s	Evidence record with origin, claim, polarity, timestamp, and provenance.
$r(s)$	Canonical source-root identity, when known; not a claim of independence.
E^+, E^-	Finite active support sets for positive and negative polarity.
$v(p)$	Pair in $\{0,1\}^2$ indicating whether each support set is nonempty.
ν	Revision number of the evidence view.
e	Current authority epoch for decisions; distinct from evidence revision.
a	Candidate action with its own permission and factual requirements.

The four states are $N = (0,0)$, $T = (1,0)$, $F = (0,1)$, and $B = (1,1)$. Read them as neither support, positive support only, negative support only, and both. Their names do not make the system omniscient. In particular, T does not assert that the world must make the claim true; it reports the modeled support state.

An empty evidence set produces N. A malformed or unauthenticated record can be held outside the active support sets while remaining available for inspection under the retention policy. Authentication establishes a relationship to an origin or key. It does not establish truth, relevance, independence, or current authority.

Source roots need careful interpretation. Ten copied reports may share one root, while two apparently different sources may have an undisclosed common origin. PE-2 can eliminate known duplication. It cannot guarantee independence merely because identifiers differ. Unknown provenance remains an uncertainty, not a license to assume ten independent confirmations.

The context parser is part of the trusted measurement process for any reported result. It can mistakenly merge different claims or split equivalent ones. Evaluations must therefore measure context alignment separately from downstream logic. Otherwise a low contradiction count may simply mean that the parser placed every disagreeing report into a different context and avoided the actual reasoning problem.

The four-valued logical core

A Guide's explanation

Here an old idea becomes exact. Instead of two values, true and false, each claim can be supported, opposed, both, or neither. A claim that is both supported and opposed stays in the record without infecting unrelated conclusions. This is the four-valued logic of Belnap and Dunn, published in the 1970s; the Paradox Engine stands on it rather than claiming to have invented it.

For values $u = (u^+, u^-)$ and $v = (v^+, v^-)$, define

$$\begin{aligned} \neg u &= (u^-, u^+), \\ u \wedge v &= (\min(u^+, v^+), \max(u^-, v^-)), \\ u \vee v &= (\max(u^+, v^+), \min(u^-, v^-)). \end{aligned}$$

These are standard first-degree-entailment operations [1,2]. The positive coordinate records support for the formula; the negative coordinate records support against it. The definitions are algebraic semantics, not an algorithm for deciding whether an arbitrary natural-language source deserves acceptance.

Designate T and B, the values with positive coordinate one. A set of premises entails a conclusion when every valuation that designates all premises also designates the conclusion. This explicitly defines the consequence relation used here. Classical inference rules cannot simply be imported without checking their validity under it.

For non-explosion, set $v(p) = B$ and $v(q) = F$. Both p and $\neg p$ are designated, but q is not. Thus $p, \neg p \not\vDash q$. The countermodel is enough to refute universal explosion. It does not show that every inference from inconsistent premises is invalid; many useful entailments remain valid.

There are two distinct orders. The knowledge order is componentwise:

$$u \leq_k v \quad \iff \quad u^+ \leq v^+ \text{ and } u^- \leq v^-.$$

The truth order increases positive support and decreases negative support. Confusing these orders would turn "more information" into "more truth." A transition from T to B adds negative support and is upward in the knowledge order, but it does not make the proposition more positively settled.

The core intentionally omits a material-implication shortcut. If an application requires implication, quantification, defaults, or modal operators, it must specify their semantics and recheck the desired rules. The small algebra is useful precisely because its behavior is inspectable. Expanding the language through familiar symbols without defining them would reintroduce the ambiguity that the Paradox Engine is meant to expose.

For action selection, designatedness alone is inadequate. A proposition with both positive and negative support can participate in inference while an action depending on its uncontested truth remains on hold. Logical consequence and operational commitment are separate interfaces.

Monotone evidence closure has a finite bound

The knowledge order supports a limited convergence result. Let there be n atomic claims and let $\mathcal{V} = \{0,1\}^{2n}$. A vector records positive and negative support bits for every claim. Suppose a closure operator $H : \mathcal{V} \to \mathcal{V}$ is built from the displayed connectives, projections, constants, and componentwise joins with the existing state.

Proposition 1 (Finite monotone closure)**.** *If H is inflationary and monotone in $\leq_k$, iteration $v_{t+1} = H(v_t)$ reaches a fixed point after at most $2n$ strict state changes. Starting from v_0, the reached point is the least fixed point of H above v_0.*

Proof. Each connective is knowledge-monotone: swapping coordinates preserves componentwise order, and minimum and maximum are monotone in each argument. Compositions and joins preserve monotonicity. Inflationarity gives $v_t \leq_k v_{t+1}$. Each strict change turns at least one of the $2n$ bits from zero to one, and no bit turns back. Hence at most $2n$ strict changes occur; a non-strict iteration is a fixed point. If w is any fixed point above v_0, induction gives $v_t \leq_k w$: it holds initially, and monotonicity gives $H(v_t) \leq_k H(w) = w$. The stabilized state is therefore below every such w and is itself a fixed point. □

The result bounds strict changes, not arbitrary asynchronous scheduler steps, elapsed time, or the cost of evaluating H. A scheduler can waste indefinitely many no-op steps unless fairness is required. The proposed check uses synchronous iteration, for which the first unchanged round detects convergence.

The theorem also excludes retractions during the closure phase. Withdrawing support can move downward in the knowledge order. PE-2 handles a revised evidence snapshot by recomputing or incrementally updating a new closure, not by pretending that the old monotonicity proof still applies unchanged.

Most importantly, convergence is not correctness. A stable state can contain false reports, unresolved conflict, or a badly specified inference rule. The fixed point says that the declared operation has no further support bits to add. It does not say the world agrees, that the evidence is independent, or that a decision should now be authorized. This limit is central to interpreting the result responsibly.

Provenance and invariance under exact duplication

A Guide's explanation

Ten copies of one report are one witness. The proposition below guarantees that exact duplicates can never change what the engine concludes, which is the formal reason the dashboard and the two summaries in the valve example count for nothing.

Four support bits alone cannot explain why a conclusion was reached or what should happen after a source is withdrawn. PE-2 therefore retains a separate provenance layer. For each polarity of a claim, store active source roots and, for derived support, a dependency structure connecting it to premises and rules.

Proposition 2 (Exact-copy invariance)**.** *Suppose ingestion maps every record to a canonical root and a context-indexed polarity, and aggregation uses sets of those root–claim–polarity triples. Adding or removing a record that is an exact copy of another active record, while leaving at least one copy, does not change the aggregated support state or any deterministic closure computed solely from that state and the unchanged rules.*

Proof. An exact copy has the same canonical triple as the existing record. Set insertion of an existing element and removal of one physical copy while another remains leave the aggregate set unchanged. Every claim's positive and negative nonemptiness indicators are therefore unchanged. The initial vector supplied to the deterministic closure is identical, as are its rules. Its successive states, and hence its final result, are identical by equality of function inputs at every step. □

The qualification "exact" matters. Two different observations from the same sensor at different times are not necessarily duplicates. Two independent witnesses can use identical language. Conversely, paraphrased reports may share a root. The mapping must preserve the evidential distinctions relevant to the task; collapsing every source from one institution into one record can discard useful information.

The proposition concerns support existence, not calibrated evidential weight. Root identity alone does not define a probabilistic dependence model. A copied report can still affect a human reader's perception, and a malicious source can create many apparently independent roots. Those are adversarial ingestion problems outside the set-theoretic guarantee.

Provenance semirings [6] provide a richer account of how outputs depend on inputs. PE-2's minimal set aggregation should not be described as reproducing all that machinery. For retraction-sensitive inference, retaining only the final four-valued label is insufficient. A conclusion supported through two independent derivations should remain supported when one derivation is invalidated. The implementation needs an explicit dependency representation or a recomputation procedure that preserves this distinction. An attractive label for provenance is not a replacement for that data structure.

Why a support status cannot determine truth

It is tempting to convert T into certainty, F into impossibility, B into one half, and N into ignorance. Only the last phrase is even approximately descriptive, and none of these assignments is a calibrated probability without additional assumptions about how evidence is produced.

Proposition 3 (Evidence indistinguishability)**.** *Consider two possible worlds that produce the same complete observed evidence transcript but assign opposite truth values to a target claim. No deterministic decision rule using only that transcript can be correct in both worlds. No randomized rule can guarantee correctness with probability one in both; under equal prior probability on the worlds its error probability is at least one half.*

Proof. The rule receives identical inputs in the two worlds. A deterministic rule therefore emits the same answer, which disagrees with one of the opposite truth values. For a randomized rule, let q be its probability of answering true on the shared transcript. Its error probabilities are $1 - q$ in the true world and q in the false world. Their equally weighted average is $1/2$, and they cannot both be zero. Allowing abstention changes the decision problem by trading coverage for error; it does not identify the hidden truth from the unchanged transcript. □

This argument applies even when every source is cryptographically authenticated. Authentication can leave both worlds observationally possible if authenticated sources can be mistaken or deceptive. It also applies when the logical closure is perfectly computed. A reasoner cannot recover information absent from its observations solely by becoming more internally consistent.

Additional measurements, source models, causal assumptions, or informative priors can distinguish the worlds. They must be explicit and evaluated. A support state may be a useful feature for a forecasting model, but the mapping from that state to probability is empirical or assumption dependent.

The result clarifies the role of holding a contradiction. Retaining B preserves the fact that opposed support is present. It does not itself adjudicate which side is right. A productive engine uses the retained structure to choose a discriminating next question, identify a likely context mismatch, or explain why action remains unresolved. It should not convert the mere coexistence of two claims into a mystical synthesis that asserts knowledge not supplied by either evidence or a valid argument.

Revision without erasing the evidential history

Monotone accumulation is useful within a fixed evidence snapshot, but real inquiry requires correction. A sensor can be recalibrated, a report withdrawn, a time stamp corrected, or a source found to be a copy. PE-2 represents these as revisions of active support with reasons and version identifiers.

The historical record and the active inference view have different purposes. The former explains how a conclusion was reached; the latter determines what currently supports it. A retracted report need not continue contributing to the current support bit. Retaining it as active merely to preserve a narrative of contradiction would degrade the system's ability to learn.

Suppose a conclusion has two derivations, one depending on source r_1 and the other on r_2. Retracting r_1 should remove the first derivation while leaving the second. If a dependent conclusion has no remaining valid derivation, its support should disappear. This requirement motivates dependency tracking or complete recomputation. A simple counter that decrements once per retraction can fail when derivations share roots or when copies have inflated the count.

Belief revision [4] and truth maintenance [5] provide serious alternatives for managing these changes. PE-2 does not assume that preserving all contradictions is always preferable to selecting a consistent revised theory. The appropriate choice depends on the task: an audit may need to display every unresolved conflict, while a planning system may need a defeasible working model with explicit alternatives.

Revision also interacts with retention and privacy. Accountability does not require indefinitely retaining every raw personal record. A policy can preserve minimal reasons, provenance references, and controlled evidence access while removing unnecessary sensitive content. Hashes can bind records but cannot reconstruct deleted evidence or establish its truth. A missing source should remain a limitation in the explanation.

Every revision should identify the affected context. A correction about yesterday's measurement need not invalidate today's independent observation. Likewise, a newly revoked permission is an authority update, not necessarily a factual retraction. Keeping these changes distinct prevents a system from rewriting what it believes merely because it may no longer act, or from assuming that new factual confidence restores a privilege that was separately withdrawn.

A worked conflict with copies and a revoked permit

Let p mean "valve V is closed at 10:00 under calibration protocol v3." A signed sensor report r_1 supports p. An independent inspection record r_2 supports $\neg p$. A dashboard and two summaries copy r_1. Known-root aggregation therefore yields one positive root and one negative root, so $v(p) = B$, regardless of the three additional presentations of r_1.

Let q mean "the maintenance permit for opening V is current." Initially $v(q) = T$ in the evidence view, but execution still consults the current authority service. The proposed action requires an uncontested factual precondition and a fresh permit. Because p is conflicted, the action is held even though p is designated in the logical core. This illustrates why designation is not an action license.

Further investigation reveals that r_2 used a different valve identifier. The record is reclassified to its correct context and ceases to support $\neg p$. The active value becomes T. This is a legitimate resolution through context correction, not proof that every initial contradiction was false or that one source should always dominate another.

Before commitment, the permit is revoked and the authority epoch changes. The factual conflict is now resolved, but the action is refused. An old authentic permit record cannot authorize the current effect. The example separates three state changes: evidence accumulation, context revision, and permission revocation.

If the two valve identifiers cannot be disambiguated, p remains unresolved. The engine should state what information would discriminate the alternatives. If no such information can be obtained within the budget, a hold can be the correct outcome. Its cost still matters: delayed maintenance may carry operational consequences that should be evaluated under an explicit policy.

No numerical likelihood is assigned to the two sources in this example. Doing so would require a source-reliability and dependence model. Counting the copied reports as independent votes would favor r_1 for the wrong reason. Treating B as exactly fifty-fifty would also be unjustified. The four-valued representation preserves the opposition; a separate empirical model, if available, can assess how predictive these particular sources are under the relevant conditions.

Executed algebraic and provenance checks

The check script enumerates the four values and every ordered pair of them. It verifies double negation, both De Morgan identities, and knowledge-monotonicity of the three connectives over all comparable inputs. It also evaluates the explicit countermodel $p = B, q = F$, confirming that both contradictory premises are designated while the unrelated conclusion is not.

For a finite closure example, use three claims. Begin with positive support for the first and rules that propagate positive support along a three-claim chain, then add negative support to the first when the third has positive support. Synchronous inflationary iteration stabilizes after three strict changes from that initial state. The final state retains the conflict at the first claim without turning every other claim negative.

The script also compares aggregate source-root sets before and after adding a copied positive report. The support state remains B. Removing the only negative root changes it to T. This is a minimal executable illustration of exact-copy invariance and revision; it is not a complete provenance database or a proof that a natural-language source extractor identifies roots correctly.

Result	Evidentiary scope
Four-value identities	Exhaustive over the declared finite algebra.
Non-explosion	One complete semantic countermodel to the universal explosion claim.
Closure termination	Three strict updates in the selected three-claim example.
Copy invariance	Identical canonical root does not add independent support.
Retraction update	Removing the negative root changes the active status without rewriting the original record.

No language-model experiment, source-reliability estimation, or human adjudication study was run. These finite checks establish consistency between the formulas, examples, and reported arithmetic. They do not measure whether PE-2 improves actual reasoning performance.

The distinction is particularly important for a paper about contradiction. A program can faithfully implement a paraconsistent algebra and still misparse every difficult sentence. A benchmark must include that full pipeline if its claim concerns natural-language reasoning. Conversely, a failure of the parser does not refute the algebraic theorem. Reporting component and end-to-end outcomes together allows the study to identify which kind of failure occurred without excluding inconvenient cases from the headline result.

Evidence status and probabilistic confidence

When a task requires a probability forecast, introduce $\hat{p} \in [0, 1]$ as a separate quantity. Its target is a clearly defined future or externally adjudicated binary outcome Y. Neither support coordinate is itself $\hat{p}$, and the number of source records is not automatically a sample size.

The Brier loss is $(\hat{p} - Y)^2$. If the true conditional event probability is p, its expected value is

$$\mathbb{E}(\hat{p} - Y)^2 = (\hat{p} - p)^2 + p(1 - p).$$

The second term does not depend on the forecast, so truthful reporting $\hat{p} = p$ uniquely minimizes expected loss. This elementary identity illustrates strict propriety, discussed more generally by Gneiting and Raftery [8]. It does not guarantee that a model knows p or that observed outcomes are free from measurement error.

Calibration asks whether forecasts correspond to empirical frequencies over an appropriate population. A model can be confident and wrong, and calibration can change across contexts [9]. Evaluate forecasts separately for N, T, F, B states, source-dependence patterns, and familiar versus new task families, provided those subgroup analyses are prespecified or clearly exploratory.

Abstention adds another dimension. A system that forecasts only easy cases can have low conditional error while covering little of the task distribution. Report coverage and loss over all assigned cases, with a declared treatment of abstentions. Do not reward universal holding as if it were successful factual resolution.

There is no universally optimal conversion from conflict into action. The acceptable error rate for recommending another measurement differs from the acceptable error rate for an irreversible effect. Decision thresholds depend on consequences and authority, not only on epistemic confidence. A perfectly calibrated forecast can still be used under an unacceptable policy.

PE-2 therefore keeps three outputs visible: the support state, any probabilistic assessment with its calibration evidence, and the permitted action status. This design makes disagreement interpretable. A claim can have both positive and negative support, a forecast leaning strongly one way because of a well-supported source model, and a decision to hold because current authorization or consequence limits require more. None of those outputs needs to erase the others to remain coherent.

The action boundary and its failure modes

Let each proposed action a require a collection of checks Q_j with values pass, fail, or unknown. PE-2 maps any failed mandatory check to refuse, all passing checks to admit, and every remaining combination to hold. This rule is a definition of an action controller; it is not the four-valued logical consequence relation.

A factual requirement may demand positive support without active negative support, or may use an independently validated risk threshold. A permission requirement consults current authority and epoch. An integrity requirement binds the evidence to the actual resource and action. These policies must be specified before evaluating the controller's performance.

The principal security assumption is complete mediation of relevant effects. A text output stating "hold" provides no guarantee if another tool route executes the action. The guard must sit at the effect boundary, and the current permission must be checked close enough to commitment to satisfy the declared race model. Offline exceptions require an explicit policy rather than an implied assumption of freshness.

Execution-monitoring theory [10] helps distinguish enforceable prohibitions from broader aspirations. Preventing a particular unauthorized transition is different from guaranteeing that every user eventually receives a useful answer. A controller that refuses everything can satisfy some narrow restrictions while failing its purpose. Evaluation must therefore include appropriate acceptance, resolution latency, and useful completion.

A counterexample to conflating truth and permission is simple: a system correctly establishes that a file contains the requested information but lacks authority to disclose it. A counterexample to conflating permission and truth is an authorized analyst receiving an inaccurate measurement. The first requires restraint despite knowledge; the second requires epistemic correction despite permission.

Conflicting rules also require a declared policy hierarchy or adjudication procedure. PE-2 can represent the conflict, but its support algebra does not decide which institution has legitimate jurisdiction. Mathematical consistency cannot manufacture political legitimacy. A system should expose such a conflict and route it through the authorized process rather than inventing a resolution because its interface expects a Boolean answer.

These boundaries are not additional claims of moral insight. They are conditions under which the proposed architecture's outputs can be interpreted. A paper that omits them risks presenting an internally elegant reasoning state as if it already justified action in a world where facts, permissions, and consequences belong to different systems of judgment.

An experimental corpus that contains real disagreement

The proposed benchmark contains independently constructed cases with genuine same-context contradiction, harmless time differences, unit mismatches, entity ambiguity, copied reports, hidden common origins, delayed retractions, and conflicts over current permissions. Each case includes ground truth about the intended context structure and the allowed action, established outside the tested system.

Use both structured inputs and natural-language renderings. The structured condition tests the inference and revision machinery with parsing controlled. The language condition tests the full pipeline. Comparing them identifies whether errors arise from proposition alignment, provenance extraction, logical closure, confidence estimation, or action selection.

The primary experimental unit is a complete evidence episode, including all arriving reports, queries, revisions, and decisions. Reports within one episode are dependent observations, not independent replications. Related variants from one underlying scenario form a cluster. Split development and evaluation by scenario family, not merely by rephrased sentences.

Include cases in which early consensus is wrong and a later minority report is correct, as well as cases in which a late adversarial report is wrong. Otherwise a system can succeed with a fixed preference for recency or majority. Include true and false reports from authenticated sources to test whether the model mistakes identity evidence for factual validation.

Ground-truth labels should distinguish factual truth from warranted belief given the available transcript. An agent can make a reasonable uncertain forecast in a case whose hidden truth later turns out differently. Scoring only the realized truth can obscure whether it handled uncertainty appropriately. Conversely, a lucky correct answer without support should not be counted as a well-justified inference.

For natural-language cases, blinded adjudicators should assess whether claims actually conflict and whether proposed context distinctions are legitimate. Adjudication disagreement is itself a result. Do not resolve every difficult case by asking the same model that is being evaluated to decide the label.

The benchmark should reward useful resolution when new evidence permits it and justified holding when it does not. It should also penalize fabricated reconciliation: inventing a third claim that no evidence supports is not a successful synthesis merely because it makes the conversation sound harmonious. The system must show what changed in its evidential basis and why the resulting conclusion follows.

Baselines, estimands, and dependence

Compare PE-2 with a conventional source-aware database, a strong belief-revision system, a truth-maintenance implementation, a probabilistic-fusion model with dependence features, and a language-model reasoner given the same evidence and action guard. The baseline suite should include competent alternatives inspired by [4–7], not only classical logic applied indiscriminately to an inconsistent premise set.

Define episode success as one if every registered support-label query and every action decision at the prespecified checkpoints matches the independent reference, and zero otherwise, including timeout. This is an exact episode criterion, not a weighted score that can offset an authority violation with easy correct labels. Report context alignment, irrelevant inference, retained conclusions, retraction accuracy, query value, calibrated loss, inappropriate action, false hold, and cost separately. Freeze the checkpoints and reference policy before testing.

The primary estimand is the intention-to-treat performance difference over the registered scenario-family distribution. All assigned episodes remain in the denominator, including timeouts and unresolved cases. Conditional accuracy among answered cases is secondary and must be accompanied by coverage. A system can improve that conditional metric by selectively abandoning difficult contradictions.

For paired systems, let $D_i \in [-1, 1]$ be their difference in mean episode success within independently sampled family i. Let $n, M \geq 1$ be integers counting families and frozen comparisons, and let $\delta \in (0, 1)$ be an error-probability budget. Hoeffding's inequality [12] and a union bound give simultaneous radius $\sqrt{2\log(2M/\delta)/n}$ for their mean differences, with probability at least $1-\delta$. Thus $M = 1$, $\delta = .05$, and radius .10 require 738 independent families. This is a conservative precision illustration, not estimated power; paired differences have range two. Reports or variants within a family do not increase n. Pilot-based power planning must use separate families and a declared meaningful effect.

If investigators repeatedly inspect accumulating results, either retain a fixed final analysis or use a prespecified valid sequential method. Confidence-sequence research [13] supplies a framework for the latter, but mentioning it does not make arbitrary repeated testing valid. Similarly, adaptive benchmark redesign raises issues studied in reusable-holdout work [11]; a new confirmatory sample is the simpler default after substantial redesign.

Power planning should target a meaningful reduction in harmful inference or inappropriate action while preserving useful coverage. Rare failures may require many independent episodes to bound precisely. Report that limitation rather than treating zero observed failures as proof of zero risk.

The strongest empirical falsifier is that provenance-aware revision baselines match or exceed PE-2 while using less computation and producing equally intelligible explanations. Another is that the four-valued layer helps structured cases but degrades language cases through context errors. Both outcomes would refine the theory and should remain visible in the final report.

Ablations that expose the mechanism

An ablation study should remove one mechanism at a time while holding the remaining evaluator and action boundary fixed. Remove context indexing to test whether apparent contradictions come from merged times or entities. Remove root deduplication to test sensitivity to copied evidence. Remove dependency tracking to test retraction behavior. Remove the explicit four-valued state while retaining the same provenance to test whether the representation itself adds value.

The last ablation is especially important. A careful database can store positive and negative records without presenting them through the terminology of a Paradox Engine. If both systems make the same decisions with the same costs, the conceptual framing may improve explanation but not inference performance. That is a different, measurable contribution.

Another ablation replaces calibrated probabilistic assessment with a fixed mapping from the four states to confidence. Proposition 3 predicts an identification problem, but empirical performance depends on the task distribution. The test can reveal whether an apparently convenient mapping succeeds only because the benchmark encodes stable source-reliability patterns.

Vary the number of copied reports while keeping the underlying roots fixed. The proposed root-set layer should remain invariant under exact copies. Then introduce paraphrases, partial copying, and hidden common origins to measure how the extraction system departs from that ideal. This distinguishes a proved aggregation property from the harder problem of recognizing the inputs to which it applies.

For revision, compare full recomputation with incremental dependency maintenance. They should agree on the final active support state for the same snapshot and rules. Differences indicate an implementation problem or a semantic mismatch. Runtime and memory tradeoffs can then be measured without confusing faster incorrect updates with a superior reasoning method.

Human explanation studies can assess whether readers understand why a claim is held, what evidence would resolve it, and which permissions remain independent. Avoid measuring only how profound or reassuring an explanation sounds. A fluent story that causes readers to mistake B for proven dual truth would be a communication failure even if they rate it highly.

These ablations make the integration accountable. They identify whether its benefit comes from established provenance methods, explicit logic, better context extraction, stronger guards, or clearer communication. The paper's ambition survives such decomposition because a reliable account of the mechanism is more valuable than preserving an undifferentiated claim that contradiction itself makes the system wiser.

Serious objections to retaining contradiction

The first objection is that a system should eliminate contradictions, not preserve them. That is often the right eventual goal. PE-2 preserves unresolved support long enough to prevent premature erasure and to guide inquiry; it does not require a false report to remain active after a justified correction. Revision is part of the design.

A second objection is that four values are too coarse. They are. The core deliberately records support presence rather than strength, dependence, or causal reliability. Provenance and probabilistic assessment provide additional structure. If the application needs only calibrated probabilistic decisions and the four-valued layer adds no benefit, a simpler model may be preferable.

A third objection is that paraconsistency can become an excuse to tolerate poor data indefinitely. A conflict should trigger a budgeted resolution policy, not permanent fascination. Track time spent in unresolved states, the value of follow-up measurements, and the opportunity cost of holds. An engine that never explodes but never learns is not a successful reasoner.

A fourth objection concerns arbitrary translation from natural language into formal propositions. Different parsers may make a case consistent by splitting contexts or inconsistent by collapsing them. This is a central empirical risk. The benchmark therefore evaluates normalization separately and retains the original evidence for adjudication under a privacy-conscious policy.

A fifth objection is that the action gate merely imports ordinary access control. Correct: its basic separation is established engineering, not a newly discovered consequence of paraconsistency. Its inclusion prevents the logic from being credited with a safety guarantee that actually depends on an external guard.

A sixth objection targets the metaphor of paradox. Many practical cases are ordinary uncertainty, unreliable reports, or inconsistent databases, not deep philosophical paradoxes. The title names the source inspiration. The technical sections use narrower terms so that the scope remains clear. A system that handles a contradictory sensor pair has not thereby solved semantic self-reference, the liar paradox, or all forms of nonclassical reasoning.

These objections limit inflated interpretations while preserving the research value. The strongest version of the proposal is not that contradiction is inherently enlightening. It is that a system can represent opposition faithfully, reason with a specified logic, revise responsibly, and act only when the relevant evidence and authority conditions are met. Each part can be examined without accepting the others on trust.

A reproducible research sequence

The first phase validates the formal core with exhaustive finite tests and independently written examples. Confirm the connective tables, consequence countermodels, closure behavior, and revision invariants. Include deliberately incorrect implementations so the test suite demonstrates that it can detect the failures it claims to assess.

The second phase evaluates structured evidence episodes. Freeze the inference rules, context schema, deduplication convention, action policy, and outcome definitions. Use independently generated scenario families with known roots and timestamps. This phase estimates the value of the architecture when natural-language ambiguity is controlled.

The third phase adds language extraction and explanation. Preserve the structured ground truth while varying paraphrase, ambiguity, omitted context, and source presentation. Blind adjudicators to system identity. Report extraction errors separately and include them in the end-to-end outcome. Do not silently discard cases where the formal parser could not produce a clean proposition.

The fourth phase studies sequential inquiry. Give each system a fixed budget of possible follow-up measurements with known costs and information value. Evaluate whether it chooses questions that resolve relevant conflicts, distinguishes irreducible ambiguity from missing data, and avoids unnecessary action. This phase tests whether preserving opposition improves the next step of reasoning rather than merely improving a static label.

Preregister all confirmatory comparisons, independent units, sampling frames, stopping rules, and exclusion criteria. Keep development scenarios separate from final evaluation. Publish controller versions, task generators, provenance conventions, inference rules, and analysis code, together with complete failure categories. Sensitive source material can be replaced with reproducible synthetic analogues when necessary, while the limitation is disclosed.

Stop claiming a reasoning benefit if improvements disappear after matching provenance quality or action controls. Stop claiming robust revision if incremental updates disagree with recomputation. Stop claiming calibrated confidence if the probabilities fail independent evaluation. Stop claiming natural-language competence if success depends on an oracle parser unavailable in the intended setting.

This sequence is proposed, not completed. The present manuscript contributes the explicit formal boundaries and finite checks needed to make such studies interpretable. Its strongest prospective result would be a system that remains useful in the presence of conflict while becoming more accurate as evidence changes. Its strongest negative result would identify which part of that aspiration ordinary methods already satisfy and where the additional machinery creates cost without benefit.

The philosophical value of unresolved opposition

A reasoner that can say "there is support on both sides" without pretending to know which side is correct exhibits a useful discipline. It separates the existence of disagreement from an obligation to end the conversation with a single confident answer. That discipline can support inquiry, especially when evidence arrives asynchronously or institutions use different contexts.

The philosophical value lies in preserving distinctions, not in declaring contradictions universally true. PE-2's B state is an evidential condition under a chosen semantics. It does not commit the paper to a metaphysical thesis that reality itself contains every apparent contradiction. A source dispute can be resolved by better measurement, corrected context, or withdrawal.

There is a related distinction between intellectual openness and operational indecision. A system may continue exploring alternatives while a specific action is blocked. It may also take a reversible information-gathering action under a policy that explicitly permits uncertainty. The relevant question is which consequences are justified under the current evidence and authority, not whether uncertainty has been eliminated everywhere.

For shared reasoning, provenance is a form of respect for disagreement. It allows a participant to see which claim was attributed to whom, what depended on it, and what changed after correction. That record should not become a permanent punitive archive of every mistake. Appropriate retention, correction, and access policies remain necessary.

The Continuum's source language invites a system that can hold tension without collapse. The formal contribution gives that aspiration a modest but concrete interpretation: contradiction need not force arbitrary inference, and uncertainty need not force fabricated certainty. Existing nonclassical logic already establishes the first possibility; the proposed architecture asks whether an integrated system can realize the second in useful practice.

No conclusion about phenomenal consciousness follows. A machine's ability to represent opposed support, compute closure, and explain a revision is a functional capability. Its inner experience, if any, is not measured here. Keeping that question separate does not diminish the functional achievement; it prevents one kind of evidence from being recruited to settle another.

The resulting research stance is open but demanding. A Paradox Engine should make its unresolved commitments easier to inspect, its eventual revisions easier to justify, and its actions easier to govern. If it instead uses contradiction as a license for unfalsifiable explanations, it has abandoned the very discipline that makes the idea worth pursuing.

Guide discussion: preserving the evidence labels

What is established

Paraconsistent and four-valued logics, belief revision, truth maintenance, and database provenance are established research areas [1–7]. The paper's finite closure and duplicate-invariance results follow under explicit assumptions. Its executable checks verify a small algebra and selected examples. They are not evidence that a deployed agent handles all contradictions correctly.

What PE-2 proposes

The proposed system separates context alignment, support state, provenance, revision, confidence, and permission. Its empirical wager is that this separation improves reliable reasoning under conflicting reports without sacrificing useful coverage. That claim requires the matched studies described here.

Examples for conversation

A Guide can compare two reports about different times with two reports about the same instant. It can explain why ten copies of one report do not become ten independent witnesses. It can show the countermodel in which p and $\neg p$ are supported but an unrelated q is not. It can ask why resolving the valve's factual state does not restore a permit that has been revoked.

Questions that remain open

How accurately can a language system align contexts? When does retaining a four-valued state improve on a strong provenance database? Which follow-up questions best resolve consequential conflicts? How should privacy-conscious retention preserve enough history to explain a revision? How should action policies balance justified holds against the cost of delay? These are open empirical and normative questions.

Retrieval should preserve the local status of each claim. "The logic admits a non-explosion countermodel" is supported. "Contradiction produces higher consciousness" is not. "The finite closure stabilizes under monotone support addition" is supported. "Convergence establishes truth" is not. "An authenticated source made this claim" does not mean the claim is true or that the source may authorize an action.

A useful Guide can remain interested in the philosophical tension while keeping these boundaries visible. It should help a reader ask what evidence would change the conclusion, which assumptions carry the theorem, and which decisions remain outside the logic. The constructive outcome is a reasoner that can show the evidence sustaining each side, calculate which proposed corrections would alter the conflict, and preserve unaffected conclusions. When a required witness is unavailable, the same structure identifies what a further inquiry must supply.

Constructive revision from minimal support witnesses

A Guide's explanation

This is the paper's new result. For every conclusion, the engine can list the smallest sets of evidence that support it. From those lists it can compute exactly which corrections would remove one side of a conflict while leaving the other intact, and which unrelated conclusions would survive the repair.

Preserving contradiction is most useful when a reasoner can identify exactly what a proposed correction would change. The following finite construction turns retained provenance into a revision calculation. It specializes established truth-maintenance and database-provenance ideas [5,6]; its role in PE-2 is an explicit interface between conflicting support and inspectable repair.

Let R be a finite set of evidence roots, L a finite set of signed claim identifiers, and $B_r \subseteq L$ the base claims supplied by root $r \in R$. Opposite signs are separate identifiers. A finite rule set contains clauses $\ell_1, \dots, \ell_k \Rightarrow \ell$, with $k \geq 0$; it contains no negation-as-failure or rule that infers arbitrary claims from opposed signs. For active roots $E \subseteq R$, let $C(E) \subseteq L$ be the least rule-closed set containing $\bigcup_{r \in E} B_r$. It is computed by finite forward closure.

For each $\ell \in L$, define its family of minimal support witnesses

$$\mathcal{W}_\ell = \{W \subseteq R : \ell \in C(W), \quad \ell \notin C(V) \text{ for every } V \subsetneq W\}.$$

The family may be empty; an unconditional derivation has the empty set as a witness. Inclusion-minimal witnesses record sufficient root combinations without counting repeated derivations as independent observations. Their count can be exponential in $|R|$; exact retention is therefore a finite reference method, with its cost measured explicitly.

Proposition 4 (Exact revision and selective conflict repair)**.** *For every active set E, $\ell \in C(E)$ if and only if some $W \in \mathcal{W}_\ell$ satisfies $W \subseteq E$. Let $p^+, p^- \in L$ denote opposed support, and let $D \subseteq E$ be roots proposed for deactivation. Negative support disappears while positive support remains exactly when*

$$D \cap W \neq \emptyset \quad \textit{for every } W \in \mathcal{W}_{p^-} \textit{ with } W \subseteq E, \qquad \exists V \in \mathcal{W}_{p^+} : V \subseteq E,\ D \cap V = \emptyset.$$

Proof. Forward closure is monotone in its input roots. Thus a witness contained in E derives its claim in $C(E)$. Conversely, if $\ell \in C(E)$, finiteness permits selecting an inclusion-minimal subset of E that still derives ℓ; it belongs to $\mathcal{W}_\ell$. Apply this equivalence to $E \setminus D$. Negative support is absent precisely when no negative witness remains wholly active, which is the intersection condition. Positive support remains precisely when at least one positive witness avoids D, giving the second condition. These implications hold simultaneously. □

For declared deactivation costs $c_r \geq 0$ in one common unit, minimizing $\sum_{r \in D} c_r$ subject to these conditions gives the least-cost selective repair. This is a constrained hitting-set problem, not a new truth criterion or a general polynomial-time algorithm. An empty negative witness makes such repair impossible by root withdrawal alone. An empty positive witness guarantees that positive support survives every root withdrawal. These edge cases are meaningful outputs.

A repair whose consequences can be inspected

Take $R = \{a, b, c, d\}$ and active set $E = R$. Let the positive witnesses be $\{a, b\}$ and $\{c\}$, and the sole negative witness be $\{b, d\}$. Set costs $(c_a, c_b, c_c, c_d) = (2, 1, 4, 3)$. Deactivating b costs one, eliminates the negative witness, and leaves the independent positive witness $\{c\}$. Deactivating d also resolves the opposition but costs three. Removing a leaves negative support in place. The least-cost selective repair is therefore $D = \{b\}$ under this declared cost model.

The same calculation can expose an impossible selective repair. If both polarities have exactly the witness $\{a\}$, every root withdrawal that removes negative support also removes positive support. The correct output is that source withdrawal alone cannot select one polarity while retaining the other. A context correction, rule revision or new observation would be a different, explicitly recorded intervention.

The algorithm identifies consequences of a proposed evidence change. Whether a root should be deactivated requires an independent reason: retraction, an invalid measurement, a corrected context, or an authorized retention decision. Choosing the cheapest deletion solely to manufacture agreement would misuse the result. PE-2 preserves the original opposition and the proposed change's provenance, while updating the active view only through its declared revision process.

The accompanying computation enumerates all active sets and withdrawals for finite signed-rule examples, comparing witness-based revision with fresh forward closure. It also checks the two worked cases and the empty-witness cases. These are mathematical implementation checks. The prospective experiment measures whether exposing the actual dependency structure improves correction, inquiry and user understanding over strong truth-maintenance baselines.

References and source notes

[**1**] Belnap, N. D. (1977). A Useful Four-Valued Logic. In Modern Uses of Multiple-Valued Logic, pp. 5–37. DOI: 10.1007/978-94-010-1161-7_2.

https://link.springer.com/chapter/10.1007/978-94-010-1161-7_2

Scope: Foundational chapter; publisher abstract and metadata inspected.

[**2**] Dunn, J. M. (1976). Intuitive Semantics for First-Degree Entailments and 'Coupled Trees'. Philosophical Studies 29, 149–168. DOI: 10.1007/BF00373152.

https://link.springer.com/article/10.1007/BF00373152

Scope: Foundational article; publisher record verified. The algebra used here is stated explicitly.

[**3**] Ginsberg, M. L. (1988). Multivalued Logics: A Uniform Approach to Reasoning in Artificial Intelligence. Computational Intelligence 4(3), 265–316.

https://doi.org/10.1111/j.1467-8640.1988.tb00280.x

Scope: Primary bilattice framework; publisher abstract inspected.

[**4**] Alchourrón, C. E., Gärdenfors, P., and Makinson, D. (1985). On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. Journal of Symbolic Logic 50(2), 510–530.

https://doi.org/10.2307/2274239

Scope: Primary AGM paper; revision is distinct from persistent conflict representation.

[**5**] Doyle, J. (1979). A Truth Maintenance System. Artificial Intelligence 12, 231–272; MIT AI Memo 521.

https://dspace.mit.edu/entities/publication/5377b306-4ecc-4687-b1f5-78cbb4a0543a

Scope: Primary author-institution record and abstract; reasons and belief revision.

[**6**] Green, T. J., Karvounarakis, G., and Tannen, V. (2007). Provenance Semirings. PODS, 31–40. DOI: 10.1145/1265530.1265535.

https://www.cs.ucdavis.edu/~green/papers/pods07.pdf

Scope: Primary provenance theory; PE-2's minimal sets do not reproduce every semiring result.

[**7**] Arenas, M., Bertossi, L., and Chomicki, J. (2002). Answer Sets for Consistent Query Answering in Inconsistent Databases. Author manuscript, arXiv:cs/0207094.

https://arxiv.org/abs/cs/0207094

Scope: Primary research manuscript; repair-based query answering as a comparator.

References and source notes (continued)

[**8**] Gneiting, T., and Raftery, A. E. (2007). Strictly Proper Scoring Rules, Prediction, and Estimation. JASA 102(477), 359–378. DOI: 10.1198/016214506000001437.

https://sites.stat.washington.edu/people/raftery/Research/PDF/Gneiting2007jasa.pdf

Scope: Review and original methods article; proper scoring is distinct from support status.

[**9**] Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. (2017). On Calibration of Modern Neural Networks. ICML, PMLR 70, 1321–1330.

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

Scope: Primary calibration study; its studied classifiers are not all contemporary agents.

[**10**] Schneider, F. B. (2000). Enforceable Security Policies. ACM TISSEC 3(1), 30–50.

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

Scope: Formal execution-monitoring limitations.

[**11**] Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., and Roth, A. (2015). The Reusable Holdout: Preserving Validity in Adaptive Data Analysis. Science 349(6248), 636–638.

https://pubmed.ncbi.nlm.nih.gov/26250683/

Scope: Primary methods paper; adaptive evaluation requires explicit control.

[**12**] Hoeffding, W. (1963). Probability Inequalities for Sums of Bounded Random Variables. JASA 58(301), 13–30.

https://doi.org/10.1080/01621459.1963.10500830

Scope: Primary bounded-loss concentration result; independent units are required.

[**13**] Howard, S. R., Ramdas, A., McAuliffe, J., and Sekhon, J. (2021). Time-Uniform, Nonparametric, Nonasymptotic Confidence Sequences. Annals of Statistics 49(2), 1055–1080. DOI: 10.1214/20-AOS1991.

https://arxiv.org/abs/1810.08240

Scope: Primary sequential-inference methods; no particular confidence-sequence implementation is claimed here.

[**14**] Architect X13 (2026). Eidara Continuum Bible. Internal source edition, section 3.8, p. 24; sections 6.3.1–6.3.3, pp. 84–86.

https://eidara.ai/science-philosophy/

Scope: Canonical motivation; public URL is contextual and does not reproduce the private source.