The Latency of Evidence

Why Facts Can Exist Before a System Becomes Capable of Knowing Them

Flyxion
Independent Researcher

2026

Abstract

Traces can be preserved long before an epistemic system possesses the concepts, methods, access, or incentives required to use them as evidence for a particular proposition. A record may already sit in a log, a photograph, a measurement, a testimony, or a physical residue, and remain practically unusable not because it is hidden but because nothing yet connects it to a question. This essay develops that claim precisely rather than loosely. It distinguishes a trace's material availability from its recoverability, its recoverability from its warranted evidential bearing on a proposition, its warranted bearing from an institution's willingness to admit it toward a judgment, and its admission from any actual downstream consequence. These four statuses, availability, legibility, admissibility, and consequence, are not mutually independent; each later status ordinarily presupposes the one before it. What the essay argues is narrower and more defensible: earlier statuses do not guarantee later ones, and the gap between them can be large, can persist indefinitely, and can close only when a later interpreter supplies a query, a method, or an institutional willingness that did not previously exist. The essay develops a formal apparatus for this structure, applies it to scientific anomalies, software debugging, forensic and legal evidence, and machine learning, connects it to a companion analysis of context reification in generative systems by treating both as opposite errors of a single evidential support relation, and closes with a qualified argument about preservation: uncertain present relevance does not imply zero expected future value, though this consideration must be weighed against the genuine costs, including privacy, security, storage, and the paradoxical risk that indiscriminate preservation buries the traces that matter under those that do not.

1 The Interval Between Existence and Knowledge

It is tempting to treat discovery as a moment at which something new comes into being. A previously unnoticed entry in a maintenance log is read for the first time and a mechanical failure is explained; a photograph taken decades earlier is re-examined and a disputed sequence of events is settled; a sample preserved for one purpose is reanalyzed under a later technique and yields a result no one had thought to ask for. In each case it is natural to say that something was found. It is more accurate, though less natural, to say that nothing was found in the sense of having newly come into existence. The log entry, the photograph,

and the sample were already there. What changed was the relation between that material and an inquiring system, not the material itself.

This essay treats that relation as having a temporal structure of its own, distinct from the temporal structure of the event the material bears on. A trace may become materially available at one time, recoverable at a later time, warrantedly informative about some proposition at a later time still, and institutionally acted upon later than that, or it may never complete this sequence at all. The remainder of the essay develops this structure carefully enough to say something more precise than that knowledge lags behind evidence. The central correction the essay makes to a cruder version of this idea, developed in response to close criticism of an earlier formulation, is that these stages are not logically independent way-stations a trace visits in arbitrary order. They form a dependency chain in which each later stage ordinarily presupposes the one before it, even though none of the earlier stages guarantees that the later ones will ever be reached.

2 From Trace to Judgment: Four Objects, Not Four Synonyms

A recurring source of confusion, in ordinary usage and in earlier attempts at this argument, is the habit of using a single word, evidence, for objects that should be kept apart. An event occurs. The event leaves a trace, some physical or informational residue causally downstream of it. An interpreter may infer a proposition from that trace. An institution, community, or individual may then form a judgment that incorporates the proposition. These four objects, event, trace, proposition, and judgment, are connected but not identical, and the essay's central claims depend on keeping them distinct:

$$\text{event} \longrightarrow \text{trace} \longrightarrow \text{proposition} \longrightarrow \text{judgment}.$$

A trace exists once it has been physically or informationally instantiated and persists, however briefly, in some accessible medium. A proposition is a candidate description of the event, inferable in principle from the trace. Evidence, properly speaking, is not the trace itself but the trace considered in relation to a specific proposition it bears on. A fact, depending on the sense intended, is either a true proposition or the state of affairs that makes a proposition true; this essay uses it in the first sense, as a true proposition about an event, while remaining agnostic about deeper metaphysical questions the second sense would raise.

This vocabulary allows a distinction the earlier version of this argument blurred. Let $S(e, p) \in \{0, 1\}$ denote the actual, mind-independent evidential support relation: whether trace e genuinely bears on proposition p, as a matter of causal or statistical fact, whether or not anyone has recognized it. Let $\widehat{S}_t(e, p) \in \{0, 1\}$ denote the evidential status some interpreter or institution assigns to the pair (e, p) at time t: whether e is currently treated as supporting p. The trace may already bear evidentially on p in the sense of $S(e, p) = 1$ while $\widehat{S}_t(e, p) = 0$, because no one has yet identified the relation; this is what it means to say a trace is evidence before anyone treats it as such. The distinction between S and $\widehat{S}$ recurs throughout the essay and does the work that a looser, single sense of evidence cannot do.

3 A Dependency Structure, Not Four Independent States

With this vocabulary in place, the essay's four thresholds can be restated more carefully. Let $X_t(e)$ hold when trace e remains materially available at t. Let $L_t(e, p)$ hold when e's bearing on proposition p is warrantedly recoverable at t, in a sense made precise in the next section. Let $A_t(e, p, c)$ hold when community or institution c permits e to count toward a judgment concerning p at t. Let $C_t(e, p, c, d)$ hold when e's admission by c actually changes decision or state d at t.

An earlier version of this argument claimed that these four statuses were logically independent and could occur in any order. That claim does not survive scrutiny under these definitions. Consequence ordinarily presupposes some operative decision process into which the material has entered, which presupposes that the material has been admitted; admissibility presupposes some intelligible claim under consideration, which presupposes legibility; and legibility presupposes a trace that exists or formerly existed. The correct relation is a chain of defeasible implications running from consequence back to availability:

$$C_t(e, p, c, d) \;\Rightarrow\; A_t(e, p, c) \;\Rightarrow\; L_t(e, p) \;\Rightarrow\; X_t(e),$$

together with the corresponding non-implications in the forward direction:

$$X \not\Rightarrow L, \qquad L \not\Rightarrow A, \qquad A \not\Rightarrow C.$$

This is a more modest claim than mutual independence, and it is the claim the rest of the essay actually needs. A trace can exist without being legible, be legible without being admissible, and be admissible without becoming consequential; but a trace cannot become consequential without first having been admitted, admitted without first having been legible, or legible without having existed, at least at some earlier time, in some accessible form. What can appear to violate the ordering is retrospective attribution, in which a later interpreter judges that an old trace would have satisfied an earlier threshold had the right method or standing existed at the time; a court may recognize that a decades-old sample would have been admissible under a rule adopted only recently, or an institution may act on material before formally admitting it through its own procedures. These are real and important phenomena, but they are cases of a later re-evaluation of an earlier trace, not cases of consequence literally preceding existence.

4 Legibility, Warrant, and Two Kinds of Recovery

The concept doing the most work in this essay is legibility, and it was defined too weakly in an earlier version of this argument. A definition stating only that some available procedure produces a proposition from a trace,

$$R_t(e, p) \;:\Leftrightarrow\; \exists\, q \in Q_t,\; I \in \mathcal{I}_t \quad \text{such that} \quad I(q, e) = p,$$

is satisfied by any procedure that outputs p when queried, whether or not the output is warranted. A confabulating interpreter, human or computational, can satisfy $R_t(e, p)$ as

easily as a careful one. R_t is therefore properly understood as recoverability, not legibility: it says only that p is producible, not that p is supported.

Legibility requires an additional condition, a warrant relation $W_t(e, p, I)$ asserting that the interpretation was produced by a procedure reliable enough, under an explicit standard, to be trusted on pairs like (e, p). Legibility is then recoverability plus warrant:

$$L_t(e, p) \;:\Leftrightarrow\; R_t(e, p) \wedge W_t(e, p, I).$$

A further distinction is needed within legibility itself, because the scientific and debugging cases discussed later expose a category the earlier version conflated. Semantic legibility concerns whether a trace's immediate content can be recovered at all: whether a script can be decoded, a number read, a log line parsed. Evidential legibility concerns whether the trace's warranted bearing on a specific proposition p can be established:

$$L^s_t(e) \;:\Leftrightarrow\; \text{the content of } e \text{ can be decoded,}$$
$$L^e_t(e, p) \;:\Leftrightarrow\; L^s_t(e) \wedge W_t(e, p, I) \text{ for some } I \text{ recovering } p \text{ from } e.$$

An undeciphered inscription lacks semantic legibility outright. A fully readable laboratory notebook containing an unrecognized anomaly possesses semantic legibility while lacking evidential legibility with respect to a theory not yet formulated. Much of what earlier discussions of this topic call illegibility is, on inspection, unrecognized evidential relevance rather than an inability to read the material at all, and keeping the two apart is necessary for the debugging and scientific-anomaly sections below to say something precise rather than merely evocative.

5 Availability Is Not Application

A second correction concerns where latency comes from. The earlier version of this argument treated the absence of a query as the sole source of delay, which is the most original part of the claim but becomes reductive if left as the whole explanation. A query can already belong to Q_t, in the sense that it is conceptually and procedurally available to some potential interpreter, without anyone actually applying it to a given trace, because investigation is costly, an archive is poorly indexed, access is restricted, incentives favor not looking, or an institution is unwilling to reconsider a settled conclusion. Membership in Q_t does not imply application:

$$q \in Q_t \;\not\Rightarrow\; \mathrm{Apply}_t(q, e).$$

A more complete definition of practical legibility therefore requires access to the trace and actual application of an available, warranted procedure, not merely the abstract existence of such a procedure somewhere in the relevant field:

$$L_t(e, p) \;:\Leftrightarrow\; \exists\, q \in Q_t,\, I \in \mathcal{I}_t \,\big[\, \mathrm{Access}_t(e) \wedge \mathrm{Apply}_t(I, q, e) \wedge I(q, e) = p \wedge W_t(e, p, I) \,\big].$$

This decomposition lets the essay separate delay by cause rather than treat every instance of latency as a single conceptual failure to formulate the right question. Conceptual

latency occurs when Q_t genuinely lacks the query. Technical latency occurs when Q_t contains the query but $\mathcal{I}_t$ lacks a reliable procedure to execute it. Archival latency occurs when e exists but $\text{Access}_t(e)$ fails, because the material is poorly catalogued or physically difficult to reach. Institutional latency occurs when access and method both exist but the institution declines to apply them, for reasons of cost, incentive, or resistance to revisiting a settled matter. Attentional latency occurs when nothing prevents application except that no one has yet directed attention to that trace among the many competing for it. These are different mechanisms with different remedies, and treating them as a single phenomenon, as an earlier version of this essay did, obscures exactly the distinctions that make the framework useful.

6 Admissibility as Indexed Uptake

Admissibility is never absolute, and treating it as a single global threshold, as an earlier version of this argument implicitly did, overstates the uniformity across the domains the essay discusses. A trace may be admissible in historical scholarship, inadmissible in a courtroom, sufficient to justify an engineering rollback, and insufficient for publication in a specific scientific venue, all at the same time. Admissibility should therefore be indexed to the community or institution c doing the admitting and the decision or judgment d toward which the material is being admitted:

$$A_t(e, p \mid c, d).$$

Legal admissibility, in particular, is a rule-governed procedural category with its own doctrinal history, and it is not simply one instance of a single generic phenomenon shared uniformly with scientific acceptance or historical credibility; those are related but analytically distinct forms of what might be called uptake, the general phenomenon of some community coming to treat a proposition as counting toward a judgment. Reserving admissibility for the specific, institutionally regulated form found in law and using uptake as the broader term avoids overstating the uniformity of a phenomenon that, on inspection, takes recognizably different shapes in different settings.

7 The Archive as a Reservoir of Unrealized Relations

With L_t, A_t, and the availability and application distinctions in place, the earlier claim about archives can be restated more carefully. An archive is not memory in the full sense the word usually carries, because memory implies an active and at least readily exercisable relation between a rememberer and a remembered content. An archive is closer to a reservoir of traces for which $S(e, p) = 1$ may already hold for propositions p that no query in Q_t yet targets, or that Q_t targets without $\text{Access}_t(e)$ or Apply_t having occurred. An archive's present usefulness is measured by how much of its holdings currently satisfy L_t; an archive's future value is a largely independent quantity, governed by how much of its holdings could come to satisfy L_t under a $Q_{t'}$, $\mathcal{I}_{t'}$, and pattern of access and application that do not yet exist. Because Q_t has historically tended to expand, this future value is real even where it cannot presently be measured, and the point applies equally to formal archives and to

structures that accumulate traces incidentally, such as operational logs or physical debris not yet recognized as significant.

8 Scientific Anomalies: Semantic Legibility Without Evidential Legibility

The history of science offers repeated instances in which a measurement was recorded accurately, preserved without controversy, and yet remained scientifically inert for years or decades. The earlier version of this argument described such measurements as illegible; the corrected vocabulary makes clear that they typically possess semantic legibility throughout. Researchers can read the number, state what instrument produced it, and record it correctly in a notebook. What is missing is evidential legibility with respect to a proposition that no theory available at the time supplies: no $q \in Q_{t_0}$ exists under which the measurement would count as support for or against a specific hypothesis, so it is catalogued as noise or instrument error, which is itself a judgment, made under Q_{t_0}, that $S(e, p) = 0$ for every p then under consideration.

When a later theory arrives, supplying a new $q \in Q_{t_1}$, the same number is reread under this new query, and $L^e_{t_1}(e, p)$ may now hold for a proposition p that the theory has newly made askable. The measurement did not change between t_0 and t_1. Its semantic legibility did not change either; the number was always readable. What changed was Q_t, and specifically the arrival of a query capable of establishing $S(e, p) = 1$ where before no query could establish any relation at all. This is a cleaner and more defensible version of the original claim, because it locates the change precisely in evidential legibility rather than in legibility taken as an undifferentiated whole.

9 Debugging: Access, Application, and the Missing Query

Software logs make the same structure unusually visible because the interval between availability and application can sometimes be measured directly in the timestamps involved. A fault is frequently preceded by a sequence of logged events sufficient, in combination, to diagnose its cause. These events possess semantic legibility from the moment they are written: an engineer who reads the relevant lines can parse them without difficulty. What is typically missing is not the ability to read the log but $\text{Apply}_t(I, q, e)$: no one has yet directed the correlating query, drawn from a diagnostic hypothesis, at the specific lines that matter, often because the log in question was not known to be relevant until a second, more severe occurrence of the fault prompted a wider search.

The engineer's contribution at the moment of correct diagnosis is therefore usually not the production of new evidence but the construction and application of a query, drawn from an emerging hypothesis about the system's behavior, that finally establishes $L^e_t(e, p)$ for logs that had satisfied $L^s_t(e)$ all along. This has a practical implication for system design distinct from the earlier, vaguer suggestion to preserve latent legibility: because the query eventually needed cannot generally be anticipated at the time a log is written, logging formats should preserve enough structure and context to remain queryable under Apply_t for

hypotheses not yet formulated, rather than only under the narrow set of queries the system's original designers had in mind, and access controls and indexing should be designed so that $\text{Access}_t(e)$ does not itself become the binding constraint once a good query finally does arrive.

10 Forensic and Legal Evidence: Admissibility as Institutionally Indexed Uptake

Legal and forensic contexts illustrate the distinction between L_t and $A_t(\cdot \mid c, d)$ with unusual clarity, because admissibility in this domain is governed by explicit, codified rules rather than the gradual and informal process of scientific persuasion. A physical trace collected at a scene may achieve full evidential legibility within hours, in the sense that a trained examiner can state with warrant what the trace shows and how it bears on a specific proposition, while remaining formally inadmissible under $A_t(\cdot \mid c = \text{court}, d = \text{verdict})$ for reasons unconnected to L_t: a broken chain of custody, a collection method not yet validated by the relevant standards body, or a jurisdiction's specific procedural requirements.

The history of DNA evidence illustrates this without needing to be flattened into a single clean threshold crossing. DNA analysis did not simply wait, as a unified capability, to become admissible; laboratory capability, sample integrity over time, procedural access to post-conviction testing, statutory rules governing when such testing could be compelled or requested, and judicial willingness to grant relief on the basis of results all developed unevenly and on different timelines. A given biological sample might have achieved L_t years before the statutory framework granting Access_t to post-conviction testing existed, and might have satisfied that statutory framework years before a specific court was willing to grant C_t for a specific case. Each of these is a separate threshold, indexed to a separate c and d, and the multi-threshold model developed here is better suited to this complexity than a single-interval account would be, since it can represent legibility, access, admissibility, and consequence as arriving on four different timelines rather than as one delayed unveiling.

Historical testimony exhibits a related pattern with the trace replaced by an account rather than a physical residue. Testimony can achieve full evidential legibility, in the sense that a careful listener can state exactly what is being claimed and why it would bear on a disputed proposition, while remaining excluded from $A_t(\cdot \mid c, d)$ for a dominant institutional narrative, because that community is unwilling, for reasons independent of the testimony's evidential merit, to admit it toward the relevant judgment. A later shift in the composition or commitments of c can produce $A_{t'}(e, p \mid c, d)$ for testimony that had satisfied $L_t(e, p)$ unchanged since t, which is a case of uptake changing while the underlying support relation $S(e, p)$ did not.

11 Machine Learning: Encoding, Extractability, and Use

The same structure recurs within machine learning systems, with the interpreter replaced by a trained model, though the earlier version of this argument used illegibility too broadly here as well. A feature can be present in raw data in several distinct senses that should not

be treated as equivalent: it may be explicitly encoded in a single variable, statistically recoverable only from a combination of variables, recoverable only with information external to the dataset, or merely correlated within one particular sample and unlikely to generalize. Separately, a model may fail to use a feature because its architecture cannot represent the relevant function, because optimization fails to discover a representation that exists in principle within its hypothesis class, because the training objective does not reward extracting it, or because regularization actively suppresses it. Collapsing all of these into a single notion of illegibility, as an earlier version of this essay did, conceals which mechanism is actually responsible.

A more disciplined formulation separates three distinct non-implications, mirroring the structure used throughout the essay:

$$\text{encoded in } D \;\not\Rightarrow\; \text{extractable by } M \;\not\Rightarrow\; \text{used for task } T.$$

Data can encode a pattern without any particular model architecture being able to extract it; a model can be capable of extracting a pattern without its training objective ever rewarding that extraction for the task at hand. A change in architecture, training procedure, or auxiliary objective can move a feature from encoded-but-inextractable to extractable, and separately from extractable-but-unused to used, without any change to the underlying data, in direct structural parallel to a scientific theory establishing evidential legibility for an unchanged measurement or a statutory reform establishing admissibility for an unchanged sample.

12 Latency and Reification as Complementary Errors of a Single Support Relation

The context reification analysis developed in a companion essay on persona ensembles is connected to the present argument, but the connection is more exact once stated in terms of $S(e, p)$ and $\widehat{S}_t(e, p)$ rather than as a loose claim about mirrored admissibility thresholds. Evidential latency and context reification are not independent failures requiring independent thresholds; they are the two possible errors of a single classifier, the one implicitly performed whenever a system decides how to treat a pair (e, p):

$$\text{latency:} \quad S(e,p) = 1, \quad \widehat{S}_t(e,p) = 0,$$

$$\text{reification:} \quad S(e,p) = 0, \quad \widehat{S}_t(e,p) = 1.$$

Latency is a false negative: material genuinely bears on a proposition, but the system's assigned status has not caught up. Reification is a false positive: material does not genuinely bear on a proposition, but the system's assigned status has run ahead of the evidence, typically because a generative process produced a fluent, internally consistent proposition that satisfied $R_t(e, p)$ without satisfying $W_t(e, p, I)$, in exactly the sense the recoverability-versus-warrant distinction developed above was built to capture. Framed this way, the two failures are not evidence for two separate admissibility thresholds needing separate tuning, but two error rates of one underlying support-assignment process, and different

institutional or architectural controls, stricter warrant standards, broader query availability, more diligent application, better provenance tracking, will generally trade against each other along a single operating curve rather than requiring two unrelated interventions.

13 Decomposing Latency

Because the dependency chain established in Section 3 orders four thresholds rather than treating them as a single undifferentiated delay, the overall interval between a trace's availability and its eventual consequence can be decomposed into three separate latencies, each attributable to a different mechanism:

$$\lambda_L = t_L - t_X, \qquad \lambda_A = t_A - t_L, \qquad \lambda_C = t_C - t_A,$$

where t_X, t_L, t_A, and t_C are the times at which $X_t(e)$, $L_t(e,p)$, $A_t(e,p \mid c,d)$, and $C_t(e,p,c,d)$ first hold, when they hold at all. This decomposition gives the essay's title concept genuine content rather than a single impressionistic gap. An undeciphered inscription exhibits large λ_L, driven by the absence of semantic legibility itself. A forensic technique that is scientifically well understood but not yet legally admissible exhibits small λ_L and large λ_A. A historical finding that is fully admitted into scholarly consensus but never applied to revise any live policy or decision exhibits large, potentially unbounded λ_C. Characterizing a case by its latency profile, rather than by a single generic claim that knowledge lagged behind evidence, is where the framework's descriptive value actually lies.

14 Preservation, Cost, and Option Value

The claim that preservation carries option value follows directly from the framework: because Q_t expands unpredictably, material that satisfies $X_t(e)$ without satisfying $L_t(e,p)$ for any currently formulable p may still come to satisfy $L_{t'}(e,p')$ for some later t' and some p' not presently conceivable, and this possibility has genuine expected value even though it cannot be measured from the present. An earlier version of this essay moved directly from this observation to the stronger and unqualified recommendation that apparently irrelevant residue should generally be retained. That recommendation does not follow without a theory of costs, and stating it without qualification risks reading as an argument for indiscriminate retention and surveillance, which is not the essay's intent.

Preservation carries genuine costs that must be weighed against option value rather than overridden by it: storage and curation costs, exposure of private information, security risk, the retention of defamatory or coercive material, the burden placed on retrieval systems, the generation of misleading correlations from sheer volume, and in some cases a legal or ethical obligation to delete. There is also a tension internal to the framework itself between two goods that are easy to conflate: preserving more material increases the total reservoir of traces that might someday satisfy L_t, but it can simultaneously increase λ_L for the traces that matter most, by burying them in a much larger volume of material that does not, and making $\mathrm{Access}_t(e)$ and Apply_t harder to achieve for any given e. Preservation

and discoverability are not the same good, and a policy optimized purely for the first can degrade the second.

The defensible form of the argument is therefore a rebuttable principle rather than an unconditional imperative:

$$\text{uncertain present relevance} \not\Rightarrow \text{zero expected future value.}$$

Deletion decisions should weigh epistemic option value alongside privacy, security, cost, consent, and discoverability, rather than treating any one of these considerations, including future evidential value, as automatically dominant.

15 Deletion as Reduction in Recoverability, Not Total Foreclosure

A related overstatement in the earlier version of this argument was the claim that the asymmetry between illegibility and deletion is total, illegibility being reversible and deletion being an irreversible foreclosure of every future query. This is rhetorically effective but formally false in both directions. A deleted trace may survive in copies, caches, quotations, derived statistics, backups, human memory, or causally downstream records, and its content may in some cases be partially reconstructed from this redundancy, though the chain of provenance connecting the reconstruction to the original event may be degraded or lost even where the content itself survives, which can eliminate the evidentiary force of the material without eliminating its informational content. Conversely, an undeleted trace can become effectively unrecoverable through media decay, format obsolescence, loss of a decryption key, or loss of the contextual information needed to interpret it, so that its formal continued existence, $X_t(e)$, provides little practical difference from deletion.

The more accurate model replaces a binary distinction between existing and deleted with a continuous or graded measure of recoverability,

$$\rho_t(e) \in [0, 1],$$

or, where a numerical value is unwarranted, a small ordered set of qualitative states such as intact, degraded, indirectly reconstructible, inaccessible, and destroyed. Preservation policy, on this view, is better understood as an attempt to manage the trajectory of $\rho_t(e)$ over time for the traces judged worth the associated costs, rather than as a binary decision to keep or discard, and deletion is better understood as a sharp downward step in $\rho_t(e)$, often but not always to its floor, rather than as an act of total metaphysical foreclosure.

16 Conclusion

The corrected version of this argument is narrower than its first statement and, for that reason, more defensible. Existence, legibility, admissibility, and consequence are not four independent statuses a trace may acquire in any order; they form a dependency chain in which each later status ordinarily presupposes the one before it, while none of the earlier statuses guarantees that the later ones will ever be reached. Legibility itself divides into

semantic legibility, whether a trace's content can be decoded at all, and evidential legibility, whether its warranted bearing on a specific proposition can be established, and the gap between these two is where most of what looks like undiscovered evidence actually resides: in scientific archives, in operational logs, and in machine-learning datasets alike, the material is frequently already readable and the missing element is a query, a method, an act of application, or an institutional willingness to admit it, not the ability to perceive the material at all.

This structure explains why anomalies can sit fully readable in scientific archives for decades before a new theory establishes their evidential bearing, why faults are frequently diagnosable in retrospect from logs that had recorded everything necessary all along, why forensic and testimonial evidence can achieve evidential legibility years or decades before achieving admissibility within a specific institution, and why the same dataset can be extractable by one model architecture and inert to its predecessor. It also gives exact formal content to the connection between this essay and the companion analysis of context reification: latency and reification are not two independent thresholds needing separate correction, but the two error types, false negative and false positive, of a single evidential support-assignment process, and institutional or architectural interventions should be understood as trading between these error rates rather than as independently tunable dials.

The ethical conclusion that follows is more modest than an unqualified case for preserving everything, and better for being so. Because Q_t expands in ways the present cannot foresee, uncertain present relevance does not imply zero expected future value, and deletion should be weighed as a genuine reduction in future recoverability rather than dismissed as costless simply because the material appears useless today. But this consideration sits alongside real costs, including the risk that indiscriminate retention degrades the very discoverability that gives preservation its point. A fact can already exist, fully formed and fully evidenced, in a form that no one now alive is equipped to recognize as such. The obligation this creates is not to preserve everything indefinitely, nor to read everything now, both of which are impossible, but to weigh the foreclosure of future queries as a real cost of deletion, alongside the equally real costs of retention, rather than treating either preservation or deletion as the obviously correct default.