

When Correction Cannot Catch the System

Temporal Adequacy of Human Authority Over Consequential Personal State

Micheal Charles Preble

Release manuscript

September 23, 2026

Abstract

Artificial intelligence governance increasingly recognizes mechanisms of correction, deletion, contestation, review, and revocation. Yet recognition of an intervention at its point of origin does not establish that the intervention remains effective once personal state has been copied, summarized, inferred from, exported to another system, or used to initiate consequential action. Persistent computational systems therefore create a distinction between correction validity and correction effectiveness: a correction may be validly accepted while superseded state continues to shape materially relevant processes elsewhere. This paper develops temporal adequacy as a governance test for that distinction. Human corrective authority over consequential personal state is temporally adequate only when an accepted corrective act, once authoritative for the relevant state, can be traced to identified materially consequential reliance paths and, within an interval proportionate to the risk of continued reliance, brings paths within the operator's governance boundary into an appropriate governed state through update, suspension, revalidation, retirement, reversal where appropriate, or documented recovery or remedy. Governing a path does not necessarily require deletion or reversal, and temporal adequacy does not require universal synchronization. Existing scholarship already addresses contestability, persistent classification, machine unlearning, authorization propagation, revocation latency, and correction duties within particular legal regimes. Machine-unlearning research examines removal of training-point influence from trained models, while privacy-policy work argues that meaningful deletion and correction may require attention to personal information used in machine-learning systems rather than stopping at the originating record. This paper does not claim those problems as new. Its contribution is narrower: it asks when an accepted correction concerning consequential personal state retains sufficient causal reach across materially consequential reliance paths to count as effective human authority.

1. The Corrected State That Still Exists Elsewhere

Current AI products provide a simple illustration of the underlying structure. Anthropic explains that when a Claude conversation expires or is deleted, memory entries generated from that conversation are not necessarily removed with it; those memory entries can instead require separate deletion. OpenAI similarly explains that saved memories are stored separately from chat history, that deleting an originating chat does not automatically delete a separate saved memory, and that removing information completely may require action across multiple sources in which the information appears. OpenAI also notes that deletion and Memory updates can take time to propagate. These examples do not establish wrongdoing, legal injury, or even material consequence. Nor do they show that every retained memory continues to influence every later output. Their significance is narrower: a valid local change to one representation of personal state need not govern every separately maintained representation derived from or associated with it (Anthropic 2026; OpenAI 2026).

A related failure of source deletion appears in machine-learning research. Izzo, Smart, Chaudhuri, and Zou define deletion from a trained model in terms of removing the effect of specified training points and develop an approximate method that avoids complete retraining for the linear and logistic models they study (Izzo et al. 2021). King and MacKinnon make the corresponding privacy-policy argument that meaningful delete-and-correct rights may need to address personal data used within machine-learning models rather than stop at the source record (King and MacKinnon 2022). These sources occupy the model-use problem; they are not the contribution claimed here.

The present paper begins with the more general governance question that follows once a corrective act has been accepted: how should we determine whether a person's authority over consequential personal state has actually become effective across the processes in which superseded state continues to matter?

2. Correction Validity and Correction Effectiveness

A correction is valid when the authoritative process accepts that prior information is no longer current, accurate, authorized, applicable, or uncontested. Validity concerns recognition of the corrective act at the point where authority over the source state is exercised. A system may therefore truthfully report that a record was amended, a memory deleted, a permission revoked, or a classification challenged.

A correction is effective when materially consequential reliance on the superseded state is also governed. That requirement does not mean that every technical descendant must disappear. It means that relevant downstream representations, decisions, systems, or institutional processes should not silently continue treating superseded personal state as though its authority were unchanged.

In this paper, consequential personal state means retained, derived, inferred, or operationalized information about a person whose use can materially affect that person's treatment, opportunity, access, obligations, classification, authority, safety, or other practically significant conditions. It does not encompass every ephemeral preference, interface setting, or non-material personalization event.

Correction validity and correction effectiveness are therefore separate questions. The first asks whether the corrective act was recognized. The second asks whether that recognized act still has causal force where the superseded state matters.

Recent work by Freiesleben, Meding, and König helps delimit the neighboring contestability problem. They define algorithmic contestability as beginning from the possibility that a decision is wrong and seeking grounds to challenge or overturn it. They contrast this with recourse, which generally treats the decision as valid and identifies changes an affected person might make to obtain a different result (Freiesleben, Meding, and König 2026). The present paper starts one step later: corrective authority has already been exercised successfully, and the question is whether its effect reaches consequential uses of the state that was corrected.

3. Four Outcomes After Correction

Treating correction as synonymous with deletion obscures the governance problem. An upstream change can require at least four different downstream responses: invalidation, revalidation, retirement, and remedy. These outcomes answer different questions and should not be collapsed into a universal command to erase or reverse.

Invalidation is forward-looking. It asks whether superseded state may continue to be relied upon for future action. A revoked instruction, corrected attribute, or disavowed classification may need to be marked or communicated in a way that prevents new consequential reliance. Revalidation asks whether a dependent result remains justified after the upstream state changes. A downstream decision may rest partly on the corrected state and partly on independent evidence, in which case automatic reversal would be as mistaken as doing nothing.

Retirement applies where historical state may legitimately remain stored while ceasing to function as current authority. Audit requirements, legal retention obligations, safety investigations, and evidentiary needs can justify preservation without continued operational reliance. Remedy addresses consequences that have already occurred. A denial, disclosure, transaction, restriction, or lost opportunity may precede the point at which correction becomes effective. Invalidation can prevent another stale use, but it cannot by itself repair an effect already produced.

The correct governance question is therefore not whether every descendant of corrected state has been rewritten. It is whether identified materially consequential reliance paths within the operator's governance boundary have entered an appropriate governed state. Operators are not required to govern independently controlled third parties they cannot see or bind. The relevant responsibility is to identify material paths within the operator's control or notification reach, and to treat remaining paths as a documented limit rather than as evidence of complete corrective effect.

4. Temporal Adequacy

A system preserves temporally adequate human authority over consequential personal state only when an accepted corrective act, once authoritative for the relevant state, can be traced to identified materially consequential reliance paths and, within an interval proportionate to the risk of continued reliance, brings paths within the operator's governance boundary into an appropriate governed state through update, suspension, revalidation, retirement, reversal where appropriate, or documented recovery or remedy.

The phrase identified materially consequential reliance paths places deliberate limits on the test. The framework does not demand omniscience. Operators may be unable to discover every transient technical copy, historical log, or independently controlled record. Technical dependency alone is also insufficient. The relevant paths are those capable of materially affecting the person concerned. For identified paths beyond the operator's governance boundary, temporal adequacy does not require impossible control; it requires a defensible account of the boundary, proportionate communication or notice where available, and an appropriate recovery, escalation, or remedy process.

Corrective latency is the interval between the point at which a corrective act becomes authoritative and the point at which a defined material reliance path reaches its appropriate governed state. It is not database-replication time. It is not a predictive model, and it does not imply that all paths must converge simultaneously.

Whether a latency is acceptable depends upon what can happen during it. A stale entertainment preference may tolerate substantial delay. A superseded identity or eligibility status used immediately before denying an essential service may not. Temporal adequacy is therefore risk-bounded rather than universally instantaneous.

5. Four Propositions

Validity-effectiveness separation: acceptance of a corrective act at the source does not establish that consequential reliance on superseded state has been governed.

Corrective-latency proposition: corrective authority becomes practically weaker when stale state can produce an unmitigated material effect before the corrective act reaches the relevant reliance point.

Dependency proposition: effective correction becomes more difficult as materially relevant reliance becomes more numerous, opaque, persistent, institutionally dispersed, or difficult to reverse.

Temporal-adequacy proposition: where superseded personal state can continue shaping material consequences, meaningful corrective authority requires risk-appropriate means of invalidation, revalidation, retirement, or remedy across identified reliance paths.

6. An Occupied Field

The surrounding field is not empty. Authorization research addresses one technically important branch of the problem. Tallam describes authorization propagation as a workflow-level property of multi-agent systems and treats transitive delegation, aggregation inference, and temporal validity as distinct authorization problems (Tallam 2026). Gallo's Proof-of-Continuity similarly examines whether authority remains causally valid across an execution lineage rather than assuming that possession of an authority artifact proves continuing authorization (Gallo 2026). Das's September 2026 individual Internet-Draft moves the issue to the effectuation boundary, asking whether a revocation that becomes authoritative before a protected act commits must control that act before it becomes externally effective (Das 2026). The draft is an individual Internet-Draft, not an endorsed IETF standard.

Parakhin supplies another technical neighbor rather than a human-governance doctrine. His capability-coherence analysis contrasts time-bounded stale-authority exposure with an execution-count revocation bound whose maximum unchecked capability uses do not scale with agent execution velocity. The protected object is capability authority, not a person's accepted correction of consequential personal state. Cache coherence therefore remains a technical analogy for one class of stale reliance, not a complete model of personal-state governance, institutional consequence, or remedy.

Human-facing scholarship occupies another part of the field. Abukar's 2026 temporal-authorship articles address persistent algorithmic classification, continuing contestability, and the ability of persons to avoid indefinite binding to prior classifications (Abukar 2026a, 2026b). Machine-unlearning work separately addresses training-point influence after source deletion, while King and MacKinnon argue in the privacy context that meaningful deletion and correction may require attention to model-level use (Izzo et al. 2021; King and MacKinnon 2022). None of those claims is treated here as unoccupied territory.

The narrower question addressed in this paper is whether an already accepted correction of consequential personal state retains causal effectiveness across the particular material reliance paths through which that state continues to operate. The paper therefore supplies a bounded governance test, not a claim that propagation, temporal validity, persistence, or contestability are newly discovered problems.

7. A Bounded Legal Analogue: The FCRA

Within the Fair Credit Reporting Act's defined consumer-reporting architecture, a furnisher's internal knowledge or source-level record status does not necessarily satisfy its statutory responsibilities when reporting to consumer reporting agencies remains inaccurate, incomplete, or materially misleading after the applicable notice and investigation triggers arise. Section 1681s-2(b)(1)(A)-(D) imposes investigation, results-reporting, and, in specified circumstances, modification, deletion, blocking, and communication duties after the statutory CRA-notice path is triggered. Regulation V's direct-dispute rule separately supplies a correction-and-notification pathway. See 15 U.S.C. § 1681s-2(b)(1)(A)-(D); 12 C.F.R. § 1022.43(e). These provisions create a bounded example in which correction is not exhausted by the source record alone; they do not create a general duty to unwind all derivatives or consequences outside the CRA channel.

Federal decisions illustrate both the reach and the limit of that principle. Saunders and Gorman illustrate that a furnisher's internal records or awareness of a bona fide dispute do not necessarily end the analysis when reporting to consumer reporting agencies remains materially misleading. Seamans further recognizes that technically accurate reporting can be materially misleading through omission and treats asserted downstream credit effects as questions of causation and damages rather than as a command to reverse later credit decisions. (*Saunders v. Branch Banking & Trust Co. of Virginia*, 526 F.3d 142 (4th Cir. 2008); *Gorman v. Wolpoff & Abramson, LLP*, 584 F.3d 1147 (9th Cir. 2009); *Seamans v. Temple University*, 744 F.3d 853 (3d Cir. 2014).)

The limiting authorities are equally important. Frazier evaluates the furnisher's duty by the information it supplies to the CRA; a later display that appears stale does not itself establish furnisher liability when the

furnished information was accurate in context. Sprague likewise keeps the privately enforceable section 1681s-2(b) route tied to the statutory CRA-to-furnisher notice chain rather than direct consumer notice alone. (*Frazier v. Dovenmuehle Mortgage, Inc.*, 72 F.4th 769 (7th Cir. 2023); *Sprague v. Salisbury Bank & Trust Co.*, 969 F.3d 95 (2d Cir. 2020).)

The FCRA therefore supplies domain-specific occupancy for the proposition that correction or knowledge at the source does not necessarily exhaust corrective responsibility where materially inaccurate or misleading information continues through a defined reporting channel. It does not establish a general duty to unwind downstream computational derivatives or completed institutional decisions. Later credit denials or unfavorable terms may be relevant damages from an actionable reporting failure; they do not become a statutory command to reverse the lender's decision.

8. Model-Level Persistence Is a Neighbor, Not the Whole Problem

Model-level persistence is an important, already studied instance of lost corrective reach. Research on machine unlearning indicates that deleting a source record need not, without further intervention, eliminate the influence of a specified training point from a trained model, and approximate deletion methods have been proposed for particular model classes (Izzo et al. 2021). Privacy-policy analysis likewise argues that meaningful deletion and correction may require treatment of model-level use rather than stopping at an originating record (King and MacKinnon 2022). The present framework does not assume that every training example remains identifiable, recoverable, or causally traceable in a deployed model.

Personal state can lose corrective reach through other pathways as well. It may be copied into a summary, retained in persistent memory, exported to another organization, transformed into a classification or score, incorporated into a workflow, or translated into an institutional judgment. Those pathways do not all call for the same response. A derived result may require revalidation; a historical record may require retirement; an independently supported conclusion may remain unchanged; and an irreversible consequence may require remedy rather than attempted erasure.

Temporal adequacy therefore does not ask whether the world has been restored to the state that would have existed had the corrected information never appeared. It asks whether materially consequential reliance has been brought under appropriate governance once authority over the personal state changes.

9. Governance Implications

Correction mechanisms should be evaluated by their causal reach rather than by their interface alone. An edit control, deletion button, dispute procedure, or revocation endpoint demonstrates that corrective authority can be recognized. It does not establish the practical effect of that authority. Evaluation must therefore ask whether materially consequential reliance remains and whether the relevant system can govern it once the status of the underlying personal state changes.

The appropriate standard is consequence-sensitive rather than architecture-specific. Low-stakes and readily reversible reliance may tolerate delay or later reconciliation. High-stakes, difficult-to-reverse, or externally consequential reliance may require substantially stronger assurance that changed authority reaches the point of action before harm becomes difficult or impossible to repair. The paper does not prescribe a single technical mechanism for achieving that result.

10. Limits

Temporal adequacy does not require universal synchronization. Some state may legitimately remain historical, some downstream conclusions may be independently justified, and some external parties may lie beyond an originating operator's practical control. A system therefore does not fail merely because every technical descendant has not been rewritten.

Nor does correction automatically require reversal. Where a corrected fact was one component of a larger assessment, revalidation may confirm the existing result. Where historical retention is legally

required, retirement may preserve the record while ending current reliance. Where consequences are irreversible, effective governance may consist of mitigation, notice, compensation, appeal, or another remedy rather than fictitious restoration.

The framework does not depend upon Operant Dyad's stronger Relationship State claims. It requires only the observable proposition that computational systems can maintain, derive, transfer, and act upon personal state over time. Relationship State therefore remains upstream and closed.

Finally, speed is not the mechanism of the problem. Persistence, dependency, transformation, and continued consequential reliance create the condition. Faster execution may enlarge exposure by allowing more consequential acts to occur before correction becomes effective, as Parakhin's authorization-specific analysis illustrates, but the same governance problem can unfold over days, months, or years. Temporal adequacy concerns whether authority reaches consequences in time, not whether computation is inherently fast.

11. Conclusion

A system can genuinely accept a correction while failing to make that correction consequential. The distinction becomes important whenever personal state survives outside its originating representation: in memories, summaries, derived classifications, reporting networks, delegated workflows, institutional processes, or other materially consequential forms of reliance.

Existing research already establishes important neighboring problems. Machine-unlearning research addresses training-point influence after source deletion; temporal-authorship scholarship addresses persistent classification and exit; authorization research addresses revocation and authority continuity; and consumer-reporting law supplies a bounded case in which obligations can extend beyond internal knowledge to communication within a defined reporting network. The contribution here is not to rename those problems. It is to provide a bounded test for when recognized human corrective authority remains practically effective across consequential personal-state reliance.

Temporal adequacy does not demand universal deletion or perfect synchronization. An accepted correction is effective when identified materially consequential reliance paths are brought into an appropriate governed state within an interval suited to their risk: invalidated where future reliance should stop, revalidated where a dependent result may remain justified, retired where history may remain without continuing as present authority, and remedied where stale reliance has already produced a consequence.

A system does not preserve meaningful corrective authority merely because it allows a person to change the source. It preserves that authority only when the change can still govern the consequences that depend upon it.

References and Authorities

Academic and technical sources

- Abukar, Ghassan. 2026a. "Temporal Authorship and the Moral Illegitimacy of Pre-Authorized AI: A New Ethical Architecture for Delegated Inquiry." *AI and Ethics* 6, article 215. Published March 25, 2026. <https://doi.org/10.1007/s43681-026-01080-9>
- Abukar, Ghassan. 2026b. "Temporal Authorship as a Moral Right: Bounded Inquiry in Persistent AI Delegation." *AI and Ethics* 6, article 249. Published April 11, 2026. <https://doi.org/10.1007/s43681-026-01093-4>
- Freiesleben, Timo, Kristof Meding, and Gunnar König. 2026. "Explainable AI Isn't Enough! Rethinking Algorithmic Contestability." arXiv:2605.16041. Submitted May 15, 2026. <https://arxiv.org/abs/2605.16041>
- Gallo, Nicola. 2026. "Proof-of-Continuity: A Temporal Model for Authority Propagation in Distributed Systems and AI Agents." CoRR abs/2607.08906. <https://doi.org/10.48550/arXiv.2607.08906>
- Izzo, Zachary, Mary Anne Smart, Kamalika Chaudhuri, and James Zou. 2021. "Approximate Data Deletion from Machine Learning Models." *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, PMLR 130:2008-2016. arXiv:2002.10077. <https://proceedings.mlr.press/v130/izzo21a.html>

- King, Jennifer, and Eli MacKinnon. 2022. "Regulating AI Through Data Privacy." January 11, 2022. <https://hai.stanford.edu/news/regulating-ai-through-data-privacy>
- Parakhin, Vladyslav. 2026. "The Bureaucracy of Speed: Structural Equivalence Between Memory Consistency Models and Multi-Agent Authorization Revocation." arXiv:2603.09875. <https://arxiv.org/abs/2603.09875>
- Tallam, Krti. 2026. "Authorization Propagation in Multi-Agent AI Systems: Identity Governance as Infrastructure." arXiv:2605.05440. <https://arxiv.org/abs/2605.05440>
- Das, Sangam. 2026. "Revoked but Still Executable: Closing the Authorization-to-Effect Gap with Finality-Bound Revocation." draft-das-finality-bound-revocation-00. Individual Internet-Draft, September 16, 2026; expires March 20, 2027. Not endorsed by the IETF. <https://datatracker.ietf.org/doc/draft-das-finality-bound-revocation/>

Product documentation

- Anthropic. 2026. "Use Claude's Chat Search and Memory to Build on Previous Context." Claude Help Center. Accessed September 23, 2026. <https://support.claude.com/en/articles/11817273-use-claude-s-chat-search-and-memory-to-build-on-previous-context>
- OpenAI. 2026. "Memory in ChatGPT." OpenAI Help Center. Accessed September 23, 2026. <https://help.openai.com/en/articles/8590148-memory-in-chatgpt>

Statutes and regulations

- Fair Credit Reporting Act, 15 U.S.C. § 1681s-2(a)(1)-(2), (b)(1), (c)-(d).
- Fair Credit Reporting Act, 15 U.S.C. § 1681i(a).
- Regulation V, 12 C.F.R. Part 1022, Subpart E, §§ 1022.42-1022.43.
- Regulation (EU) 2016/679 (General Data Protection Regulation), arts. 16 and 19 (neighboring, non-load-bearing authority).

Cases

- Saunders v. Branch Banking & Trust Co. of Virginia, 526 F.3d 142 (4th Cir. 2008).
- Gorman v. Wolpoff & Abramson, LLP, 584 F.3d 1147 (9th Cir. 2009).
- Seamans v. Temple University, 744 F.3d 853 (3d Cir. 2014).
- Frazier v. Dovenmuehle Mortgage, Inc., 72 F.4th 769 (7th Cir. 2023).
- Sprague v. Salisbury Bank & Trust Co., 969 F.3d 95 (2d Cir. 2020).