

Human Agency Under Consequential State

A Policy Framework for Correction, Revalidation, and Recoverable Consequences

Micheal Charles Preble

Independent Researcher

Executive Summary

Persistent computational systems can preserve the effects of earlier classifications, scores, inferences, and decisions after the state that supported them has changed. This paper argues that correction is incomplete when the source record changes but materially dependent consequences remain operative. It distinguishes correction from revalidation and develops a Consequence Revalidation Framework for determining whether continuing consequences retain adequate authority.

The paper introduces corrective continuity: the capacity of a valid correction to remain effective across the decision and consequence pathways through which superseded state would otherwise continue to affect a person over time.

Keywords artificial intelligence | human agency | consequential state | revalidation | corrective continuity | provenance | effective remedy

Key Takeaways

Correction changes state.
Revalidation determines whether effects produced under the former state retain authority.

Corrective continuity requires valid corrections to travel across materially dependent decision and consequence pathways.

The operational lifecycle is:
Correction -> Dependency ->
Provenance -> Participation ->
Revalidation -> Remedy -> Closure.

U.S., EU, and international frameworks contain pieces of this architecture, but the cross-domain integration remains underdeveloped.

I. Introduction - The Problem After Correction

Persistent computational systems create a temporal governance problem: state about a person can continue shaping institutional treatment after the grounds supporting that state have changed.

Computational systems increasingly create classifications, scores, inferences, records, rankings, and other forms of state that affect how people are treated by institutions. That state may influence eligibility, employment, access, restrictions, financial decisions, risk assessments, prioritization, and other consequential judgments. Increasingly, it may also persist.

Persistence creates a problem that conventional correction rights do not fully resolve. Suppose information about a person contributes materially to an adverse decision. The information is later corrected. Or an inference is withdrawn. A classification is invalidated. Evidence is superseded. Authority to rely on a prior state expires.

The originating record may now be accurate. The consequence may nevertheless remain.

A restriction can continue. An eligibility determination can remain unchanged. A downstream organization can retain a decision made using the former state. A derivative classification can persist in another workflow. A past determination can become an input into a later one.

The problem is therefore not simply whether computational state can be corrected. It is whether the authority of consequences generated from that state remains revisable after the state itself changes.

Existing law recognizes many components of this problem. Data-protection regimes provide correction and, in some circumstances, downstream notification. Administrative regimes provide reconsideration. Particular statutory systems provide corrected-record review and restorative relief. Procedural doctrines protect opportunities to understand and contest consequential decisions. Human-rights frameworks recognize effective remedy, restitution, participation, and cessation of continuing harms.

But those mechanisms generally arise in different doctrinal settings and answer different questions. The unresolved governance problem is narrower:

What should happen when consequential computational state is corrected, withdrawn, superseded, or invalidated after decisions materially dependent on that state have already produced continuing effects?

This paper argues that correction alone is insufficient. A correction determines what state should now be regarded as valid or operative. It does not necessarily determine whether decisions made under the former state remain justified. Nor does it determine whether consequences produced by those decisions should continue.

That requires a separate institutional act: revalidation.

Correction changes state. Revalidation determines whether effects produced from the former state retain authority.

This paper therefore develops a Consequence Revalidation Framework for persistent computational systems. The framework proceeds through seven linked stages:

Correction -> Dependency -> Provenance -> Participation -> Revalidation -> Remedy -> Closure.

The framework does not require automatic reversal when state changes. A decision may remain valid on independent grounds. A corrected input may have been incidental. Finality, reliance, third-party authority, elapsed time, and remedial impossibility may constrain what can be done.

But automatic continuation is equally inadequate. Where materially consequential state has changed, an existing consequence should not retain authority merely because it has already entered an institutional workflow.

The central rule is therefore: correction should trigger revalidation, not automatic reversal.

This paper describes the temporal property necessary to make that rule meaningful as corrective continuity: the capacity of a material correction to remain effective across the decision and consequence pathways through which superseded state would otherwise continue to affect a person over time.

The unit of governance therefore cannot always be the model, the record, or even the individual output. It must sometimes be the continuing authority of consequential state across time.

The resulting proposition is modest but consequential: when materially consequential computational state changes, the responsible originator or recipient should determine whether materially dependent continuing consequences still possess an adequate basis.

That is the problem after correction.

II. From Informational Correction to Consequential-State Governance

The problem begins with a distinction among different kinds of state. A computational system may hold information about a person. That information may influence a decision. The decision may produce a consequence. The consequence may then persist independently of the information that originally helped produce it.

Correction law often focuses on the first layer. Consequence revalidation requires distinguishing all four.

A. Information State

Information State includes factual records, derived attributes, inferences, classifications, scores, annotations, and other representations about a person that may influence institutional action. A correction at this layer asks what should now be treated as accurate, valid, current, or operative.

Information State therefore includes more than conventional data accuracy. A record may remain historically accurate while losing authority for future use. A classification can truthfully describe what a system concluded at an earlier time and nevertheless cease to be an appropriate basis for present treatment.

B. Decision State

Decision State is the institutional judgment produced using Information State together with other information, rules, policies, and human judgment.

A simplified decision may be represented as $A + B + C \rightarrow D$. If B is later corrected, the correction does not itself reveal what should happen to D. D may have depended on B entirely; B may have contributed materially but not been necessary; B may have been present but irrelevant; A and C may independently support the same result; or the institution may no longer possess enough evidence to determine B's role.

Changing an input is not the same institutional act as reconsidering a judgment produced with that input.

C. Consequence State

Consequence State consists of the material effects produced by a decision. Those effects may include restrictions, loss of opportunity, eligibility, ranking, priority, financial treatment, employment consequences, referral, account limitation, institutional status, or other changes to the person's reachable options.

The significance of Consequence State is that it can persist after both the computation and the decision event have ended. A score can disappear while a denial remains. A model can be retired while a classification generated by it remains in another system. A record can be corrected while a restriction created under the former record continues.

D. Closure State

Closure State records whether the institution has resolved the continuing authority of consequences materially connected to the corrected state.

Closure is the documented resolution of the continuing authority of materially dependent consequences.

The four layers can therefore be represented as:

**Information State $\rightarrow$ Decision State $\rightarrow$
Consequence State $\rightarrow$ Closure State.**

E. The Boundary: Material Dependency

Not every change in Information State requires reopening every later event. The relevant boundary is material dependency. The inquiry is not whether corrected state appeared somewhere in the institutional history. It is whether

that state materially supported a decision or continuing consequence.

The relevant object is not every downstream event that followed chronologically. It is every continuing condition in which superseded state remains materially operative.

F. Consequential and Corrective Continuity

Persistent systems introduce an additional temporal dimension. Prior state can continue influencing later decisions even after the originating interaction has ended. A classification can feed another screening process. A prior denial can affect later eligibility. A restriction can alter the level of future scrutiny. Records can be copied, inherited, or transformed into derivative state.

Consequential continuity is the persistence of material effects produced by prior consequential state after the original state has entered later decisions, records, or institutional pathways.

Corrective continuity is the capacity of a material correction to remain effective across the decision and consequence pathways through which superseded state would otherwise continue to affect a person over time.

Consequential continuity describes how prior state continues to matter. Corrective continuity asks whether a valid change to that state can continue to matter as well.

Corrective continuity does not require deletion of history or reversal of every prior outcome. It requires correction to reach the locations where the superseded state still claims operative authority. A persistent system must therefore be capable of both remembering and superseding.

III. Existing Law Approaches the Problem From Different Directions

Existing law already recognizes many of the individual mechanisms necessary for consequence revalidation. Different regimes provide correction, propagation, explanation, reconsideration, participation, and remedy. What remains fragmented is the relationship among them once corrected state has already produced continuing effects.

A. European AI Governance: System Lifecycle and Corrective Action

The European Union Artificial Intelligence Act establishes lifecycle obligations for certain AI systems. Article 20 requires providers of nonconforming high-risk AI systems to take corrective action and inform specified downstream actors, while Article 72 requires documented post-market monitoring across the high-risk system's lifetime.¹

These mechanisms reflect an important shift from one-time conformity toward continuing governance. But their principal object remains the system.

A lifecycle for the system is not necessarily a lifecycle for the consequences produced through the system.

A system can be corrected or withdrawn while prior classifications, restrictions, eligibility determinations, or downstream decisions remain operative. The AI Act therefore supplies important lifecycle machinery without itself creating a general cross-domain consequence-revalidation rule.

B. GDPR: Rectification Can Propagate Beyond the Original Record

The General Data Protection Regulation comes closer to corrective continuity at the information layer. Article 16 creates a right to rectification, Article 18 permits restriction of processing in specified circumstances, and Article 19 requires communication of covered rectifications, erasures, and restrictions to recipients, subject to impossibility or disproportionate effort.²

That is significant because it establishes that correction is not always complete when the source record changes.

But informational propagation and consequence revalidation are still distinct.

Information A -> Decision B -> Consequence C.

Communicating corrected A to the recipient does not itself determine whether B must be reconsidered or whether C may continue.

Correction propagation is not necessarily consequence propagation.

C. Distributed Automated Decisions

Recent Court of Justice case law illustrates the distributed-decision problem. In *SCHUFA Holding (Scoring)*, C-634/21, the Court addressed credit scoring performed by one entity and heavily relied upon by another in consequential decisions, treating the scoring activity as capable of falling within Article 22's automated-decision framework in the circumstances before it.³

The structural point is important: consequential computational processing can be distributed across actors.

State producer -> score or inference -> downstream decisionmaker -> consequence.

In *Dun & Bradstreet Austria*, C-203/22, the Court interpreted Article 15(1)(h) of the GDPR as requiring the controller to provide meaningful information about the logic involved in automated decision-making, including profiling, in a form capable of enabling the data subject to understand and challenge the decision-making process.⁴

Those cases do not establish the framework proposed here. They demonstrate why it may be necessary: consequential computational state can travel across organizational boundaries, while meaningful contestation depends on understanding material dependency.

D. U.S. Privacy Law: Record Correction and Decision Review Are Different Acts

U.S. federal privacy law demonstrates the separation between record correction and substantive reconsideration.

The Privacy Act provides mechanisms for amendment of federal records, but courts have distinguished factual or historical correction from attempts to rewrite or reopen discretionary judgments. In *Lee v. Geren*, the district court treated amendment of factual or historical errors as distinct from using the Privacy Act to revisit substantive personnel judgments.⁵

The significance is structural. The law can recognize the state being corrected without automatically reopening the judgment produced from that state.

E. A Domain-Specific Model: Special Selection Boards

One U.S. statutory regime comes particularly close to consequence revalidation. Under 10 U.S.C. § 628, where military promotion consideration involved specified forms of material unfairness or error, a special selection board may reconsider the officer using the corrected record as it would have appeared to the original board.⁶

The structure matters more than the domain:

correction -> reconstructed decision setting -> reconsideration under corrected state.

Where the statutory process results in promotion, the regime can restore downstream effects including rank and pay. *Roth v. United States* illustrates the distinction between identifying error and obtaining the restorative consequence; corrected-state reconsideration intervenes before the ultimate relief.⁷

That regime demonstrates in concrete form the principle developed here: correction should trigger revalidation, not automatic reversal.

Its significance is conceptual. It is a domain-specific example of corrective continuity, not a general rule for computational systems.

F. Administrative Reconsideration and Finality

Administrative law contributes another piece. Agencies can possess authority to reconsider prior decisions in order to correct legal, factual, or procedural error, but that authority is bounded by statutory limits, timing, reliance, and finality. *Voyageur Outward Bound School v. United States* illustrates that reconsideration authority and its temporal limits in an agency setting; it is not a person-facing analog and is cited only for that narrower point.⁸

The tension is important. Error correction matters. Finality matters. Corrective continuity is bounded continuity.

G. International Human-Rights Frameworks: Remedy Beyond the Source Error

International human-rights materials contribute a broader remedial principle. Human Rights Committee General Comment No. 31 explains that cessation of an ongoing violation is an essential element of effective remedy, and that reparation can include restitution, rehabilitation, satisfaction, compensation, and guarantees of non-repetition.⁹

The relevance to computational governance is limited but important. Not every erroneous computational consequence

is a human-rights violation. The narrower point is that established remedial doctrine does not assume that correcting the source of an injury necessarily resolves its continuing effects.

The United Nations Guiding Principles on Business and Human Rights similarly emphasize remediation and effective grievance mechanisms characterized by legitimacy, accessibility, predictability, equity, transparency, rights compatibility, and, for operational-level mechanisms, engagement and dialogue.¹⁰

H. UNESCO and Human Agency Across the AI Lifecycle

UNESCO's Recommendation on the Ethics of Artificial Intelligence adds a human-agency and lifecycle dimension. In the paragraphs cited here, it treats transparency and explainability as preconditions for challenging AI-informed decisions and connects those requirements to responsibility and accountability across the AI lifecycle.¹¹

A person has not necessarily regained meaningful agency merely because a system's internal representation has changed. The practical question is whether the correction can affect the environment through which the person continues to be treated.

I. The Integration Gap

Across these regimes, the component principles already exist: correction, propagation, explanation, reconsideration, participation, restoration, finality, and effective remedy.

The remaining issue is how those mechanisms fit together when computational state materially changes after it has already produced continuing consequences.

This survey does not identify a general cross-domain rule requiring an institution to identify still-operative decisions materially dependent on superseded computational state; determine whether those decisions remain independently valid; address continuing consequences that no longer retain an adequate basis; provide meaningful participation where it could affect that determination; and preserve evidence showing how the matter was resolved.

The gap is therefore not simply an absence of correction rights. It is an integration problem.

The proposed framework connects existing mechanisms around a single object: the continuing authority of consequences produced from superseded computational state.

IV. Correction Is Not Revalidation

The distinction between correction and revalidation is the conceptual center of this paper.

Correction determines what information, inference, classification, or other consequential state should now be treated as valid or operative.

Revalidation determines whether a decision or consequence produced under the former state retains adequate justification after that state changes.

Correction looks at the state. Revalidation looks at what the state already did.

A. Why Correction Alone Is Insufficient

Consider $A \rightarrow B \rightarrow C$. A is consequential computational state. B is a decision materially informed by A. C is a continuing consequence produced by B. A is later corrected to A'.

That establishes that A should no longer govern future processing in the same way. It does not by itself establish whether B remains valid or whether C may continue.

Correction therefore does not answer the consequence question. Revalidation does.

B. Why Automatic Continuation Fails

Where a material premise of a continuing outcome has been withdrawn or invalidated, the mere fact that the decision was previously made does not establish that the outcome remains justified.

Historical existence is not the same as continuing authority. A consequence should not continue merely because it has already been produced.

C. Why Automatic Reversal Also Fails

The opposite rule would reverse every decision touched by corrected state. That is equally unsatisfactory.

Suppose $A + B + C \rightarrow D$. If A is invalidated, B and C may still independently support D.

The correct inquiry is: does the decision remain adequately justified without the superseded state?

Correction should trigger revalidation, not automatic reversal.

D. Four Forms of Dependency

Material dependency can be classified into four categories.

Necessary dependency: the consequence could not have been sustained without the superseded state.

Contributory dependency: the state materially affected the decision but was not its sole basis.

Incidental presence: the state was present but did not materially affect the outcome.

Indeterminate dependency: available evidence is insufficient to determine its role.

These categories prevent minor corrections from triggering indiscriminate reopening and prevent uncertainty from being treated as proof that the corrected state did not matter.

E. Historical Basis and Current Basis Must Remain Separate

Suppose a restriction was imposed at T1 because of invalid State A. At T2, A is corrected. At T3, the institution discovers valid State B, which independently supports continuation.

B may justify the restriction prospectively. But B did not justify it at T1.

Current justification cannot rewrite historical dependency.

That distinction becomes particularly important when determining retrospective remedy.

F. The Stopping Rule

Revalidation does not require reconstructing every event that followed from corrected state.

Trace consequences until the corrected state is no longer a material basis for an operative decision, status, restriction, entitlement, ranking, classification, obligation, or other consequential condition.

This separates causal history from operative dependency. The framework is concerned with the latter.

G. Finality Is Part of Revalidation

Finality and corrective continuity are not opposites. A mature framework must consider elapsed time, severity of the consequence, reliability of available evidence, independent grounds, third-party reliance, statutory limits, and feasibility of restoration.

Those considerations may limit reopening or remedy. They do not eliminate the conceptual distinction between a consequence that remains because it was substantively revalidated and one that remains because governing law no longer permits reopening.

H. The Temporal Governance Problem

Prior state can become embedded in later institutional pathways. The source record can change while the person's trajectory remains unchanged.

Correction changes the state. Dependency identifies where the old state mattered. Provenance establishes whether that relationship can be reconstructed. Participation permits meaningful challenge where it could affect the result.

Revalidation determines continuing authority. Remedy addresses consequences that no longer retain an adequate basis. Closure records the result.

V. The Consequence Revalidation Framework

The framework that follows is a proposed institutional architecture. It synthesizes existing legal mechanisms and analogues but does not claim that existing law presently imposes each element as a general cross-domain duty.

When consequential computational state materially changes, the responsible actor should determine whether materially

dependent continuing consequences still possess an adequate basis.

The canonical lifecycle is:

Correction -> Dependency -> Provenance -> Participation -> Revalidation -> Remedy -> Closure.

Trigger is an entry condition to Dependency, not a separate pipeline. Propagation is part of Remedy and Closure: it carries a revalidation result to locations where superseded state remains operative. The five operational groupings below implement that single lifecycle; they do not replace it. Provenance and Participation are developed in Sections VI and VII.

A. Trigger

Revalidation is triggered by a material correction, withdrawal, supersession, or invalidation of consequential computational state where there is a still-operative consequence and a plausible path of material dependency between the former state and that consequence.

A trivial clerical correction with no plausible connection to an operative consequence need not proceed further. The trigger follows live operative dependency, not every event that followed chronologically.

B. Dependency

Dependency is necessary, contributory, incidental, or indeterminate.

Where provenance is inadequate, the indeterminate category should remain visible. Missing provenance must not silently become "probably did not matter."

C. Revalidation

Where material dependency exists, or cannot reasonably be excluded because the record is inadequate, the institution should determine whether the outcome remains justified under corrected state.

Possible results include revalidated unchanged; revalidated with modification; suspended pending review; reversed; prospectively remediated; restoratively remediated; or unresolved.

The central inquiry is: would the decision remain adequately justified under the corrected state?

D. Propagation Within Remedy

Revalidation is incomplete if the result remains confined to the originating system while superseded state continues to operate elsewhere. Propagation belongs to remedy: it is the movement of correction and revalidation results to the places where a consequence still claims authority.

Trace consequences until the corrected state is no longer a material basis for an operative decision, status, restriction, entitlement, ranking, classification, obligation, or other consequential condition.

The rule follows continuing material dependency, not every conceivable causal effect.

E. Closure

Closure requires more than recording that a correction. The closure authority should be able to show what changed; what decisions were reviewed; what consequences remained operative; how dependency was assessed; whether the outcome remained valid; what remedy occurred; whether the downstream propagation was required; who possessed authority to resolve the matter; and when closure occurred.

Closure marks the point at which the relevant consequence path has been resolved.

F. Allocating Responsibility

Roles should be separated where authority is separated. The originator corrects source state, identifies known dependency paths, and transmits correction or revalidation information. The recipient or deployer determines whether its own continuing decision or consequence remains justified. A closure authority records whether the matter was actually resolved. Propagation is not control: an originator may owe notice and transmission without possessing authority to reverse a recipient's independent decision.

Its object is not the model. Its object is the continuing authority of the consequence.

VI. The Provenance Problem

The framework depends on being able to answer a difficult question: what role did the superseded state actually play?

A final outcome may reflect source data, derived attributes, model outputs, business rules, prior classifications, human judgment, overrides, and downstream processing. Some of those components may not be retained in a form that permits later reconstruction.

If uncertainty always favors continuation, poor provenance becomes a shield against accountability. The less reconstructable the process, the harder its consequences become to challenge.

A. The Provenance Deficit Rule

Where an affected person identifies a corrected consequential state and a plausible materially dependent consequence, the responsible actor should bear a burden of production sufficient to demonstrate whether that consequence remains independently valid. If materially necessary provenance is unavailable because the responsible actor did not preserve it, the resulting uncertainty should trigger proportionate revalidation measures rather than be treated as evidence that the consequence remains valid.

The rule does not presume invalidity. It assigns responsibility for producing the evidence needed to determine dependency.

B. Production Duty

The affected person should first identify a material correction, a concrete continuing consequence, and a plausible path of material dependency.

Once that threshold is met, the responsible actor should produce enough evidence to permit a reasonable dependency assessment, such as source information, material derived attributes, model or rule versions, decision records, human-review notes, overrides, timestamps, and downstream transfers.

The actor that continues to enforce the consequence should be able to identify what materially supported it.

C. Revalidation Under Incomplete Provenance

If provenance shows material dependency, revalidation follows. If it shows the state was incidental, the inquiry may. Where dependency cannot reasonably be excluded because relevant provenance is missing, the responsible actor should not default to continuation without further inquiry.

If dependency is established or cannot reasonably be excluded because provenance is inadequate, the consequence should be revalidated against corrected state.

D. Protective Response

Where provenance remains inadequate and the consequence is serious or continuing, proportionate protective measures may include a fresh determination, temporary suspension, prospective remediation, or explicit classification of the matter as unresolved.

Uncertainty created by missing provenance should be governed, not ignored.

E. Provenance States

For operational purposes, provenance can be classified as reconstructable, partially reconstructable, or non-reconstructable.

Partial reconstructability should remain visible as uncertainty rather than being converted into artificial certainty.

F. The Core Evidentiary Rule

A responsible actor should not be able to convert failure of reconstructability into evidence of continuing validity.

If missing provenance increases revalidation burden, institutions have reason to preserve the evidence necessary to justify continuing consequences.

G. Provenance Across Time

A decision may be reconstructable when made and become non-reconstructable years later because logs expire, systems migrate, models are retired, or human records disappear.

The duration of a consequential effect should inform the duration for which material provenance remains available to justify it.

This is not a demand for universal permanent retention. It is a governance relationship between the persistence of a consequence and the persistence of the evidence needed to justify it.

H. Provenance as Correction Infrastructure

Correction identifies what changed. Provenance identifies where it mattered. Revalidation determines what remains valid. Propagation carries the result forward. Closure records the resolution.

Without provenance, corrective continuity can break between correction and consequence.

VII. Participation in Revalidation

Within this proposed architecture:

Revalidation should not always be an internal process conducted entirely outside the affected person's view.

Where participation could materially affect the result, the person should have a meaningful opportunity to understand and contest the basis for continuation.

Participation should be sufficient for the affected person to understand and materially contest the basis for continuing a consequential outcome, but should scale with severity, uncertainty, reversibility, and the practical value of participation.

A. Notice

The person should ordinarily receive notice that relevant state changed, a continuing consequence is under review, and an institution is responsible for the revalidation.

B. Basis Disclosure

Participation requires enough information to understand why the consequence may continue: what consequence remains, what state changed, whether the institution believes the old state materially affected the outcome, what independent grounds now support continuation, and what issues remain open to challenge.

The standard is not maximum disclosure. It is meaningful contestability.

C. Meaningful Response

The affected person should have a practical opportunity to challenge dependency classification, factual errors, omitted information, claimed independent grounds, unidentified downstream effects, or the scope of the consequence itself.

Participation can reveal information that internal provenance misses.

D. Reasoned Determination

A bare statement such as "the matter was reviewed and remains appropriate" is insufficient where the consequence is material and the basis for continuation is genuinely contestable.

The determination should identify whether the outcome was revalidated unchanged, modified, reversed, suspended, remediated, or unresolved.

E. Review Path

Where further review exists, the person should be able to identify it. The framework does not create a universal appellate mechanism. It requires visibility into the mechanism that actually exists.

F. When Participation Adds No Value

Not every case warrants extensive participation. If a correction triggers a purely mechanical recalculation and no discretionary issue remains, elaborate procedure may serve little purpose.

Participation is required where participation could materially affect revalidation.

G. Emergency Conditions and Declined Participation

In limited settings, immediate removal of a consequence may create substantial risk. Participation may then be delayed, but it should not disappear. Emergency continuation should be temporary, documented, and followed by expedited review.

An affected person may also decline to participate. Revalidation may proceed on the available record, with closure noting that participation was offered and declined.

H. Participation and Corrective Continuity

Participation is especially important in distributed systems because the affected person may be the only actor who sees the entire continuity problem. One system sees the record, another sees the decision, another sees the restriction. The person experiences the combined effect.

Participation therefore serves both fairness and discovery.

VIII. Remedy and Closure

Within this proposed architecture:

Revalidation determines whether a consequence remains valid. Remedy determines what should happen when it does not. Closure determines when the institution has completed that process.

Remedy should address the continuing or accumulated effects of invalid consequential state to the extent legally and practically available, while preserving distinctions between

historical invalidity, current justification, third-party autonomy, and remedial impossibility.

A. Propagation Is Not Control

The institution responsible for correcting state may not control every consequence produced from it. It may nevertheless have responsibilities to communicate the correction, identify the state being superseded, transmit relevant revalidation results, and document propagation.

Propagation is not control.

B. Historical Invalidity and Current Justification

A later-valid basis may justify prospective continuation, but it cannot rewrite the historical basis of the original consequence.

Past remedial questions and present authority should therefore remain distinct.

C. Prospective and Restorative Remedy

Prospective remedy may include lifting a restriction, restoring eligibility, changing a classification, ending reliance on superseded state, modifying a ranking, or replacing the prior decision with a fresh determination.

Where law and practical capacity permit, restorative measures may include reinstatement, retroactive entitlement, restored priority, financial restoration, correction of derivative records, or return to the position the person would otherwise have occupied.

Where the consequence can be reconstructed, corrective continuity may sometimes require reconstructing the person's position as well.

D. Substitute Remedy

Some consequences cannot be undone. Opportunities expire, time passes, third parties act, institutions change, and evidence disappears.

Where literal restoration is impossible, substitute relief may be appropriate where available. Remedial impossibility should change the form of remedy, not erase the remedial question.

E. Third-Party Consequences

Where third parties control downstream consequences, responsibilities should be separated. The originator may need to correct, propagate, identify recipients, and preserve evidence of transmission. The recipient's obligations depend on its own authority and governing law.

F. Emergency Continuation

Where immediate termination would create substantial risk, temporary continuation may sometimes be justified. It should require identified authority, documented risk, a reason, a defined duration, a review deadline, and expedited revalidation.

Emergency continuation must be temporary, justified, and reviewable.

G. Closure States

Possible closure states include: revalidated unchanged; revalidated with modification; reversed; prospectively remediated; restoratively remediated; substitute remedy applied; temporarily continued under emergency exception; and unresolved.

H. Closure Documentation

The closure record should identify what state changed, how dependency was classified, the provenance status, whether participation occurred, the revalidation result, the remedy, whether propagation occurred, the authority for the disposition, and the closure date.

The closure record resolves the continuing authority of materially dependent consequences; it is more than a timestamp showing that correction occurred.

IX. Hard Cases and Limiting Principles

A consequence-revalidation framework should be judged by how it handles difficult cases, not obvious ones.

A. Independent Grounds Support the Same Outcome

If corrected State A materially contributed to Decision D but independent State B would have supported the same result, the appropriate disposition may be revalidated unchanged. The closure record should nevertheless identify B as the independent basis rather than rewriting history as though A never mattered.

B. Necessary or Contributory State

If the superseded state was necessary and no independent basis remains, the consequence requires either a new basis or a new disposition.

Where several factors supported the original decision, the question is whether the remaining factors still adequately support the outcome once the corrected factor is removed.

C. Provenance Is Inadequate

Where the role of corrected state cannot be reconstructed, the institution should not conclude that inability to prove materiality establishes validity. A concrete consequence plus a plausible path of material dependency triggers proportionate revalidation under the Provenance Deficit Rule.

D. Speculative Connections

The framework does not require institutions to investigate every possible event following a correction. An affected person should identify a material correction, a concrete consequence, and a plausible path of material dependency.

E. Long Dependency Chains

Revalidation should extend as far as material operative dependency extends, and no farther. The inquiry ends once the superseded state is no longer a material basis for an operative condition.

F. Irreversible Consequences

Where literal restoration is impossible, the inquiry shifts to substitute remedy where available. Irreversibility changes the remedial form; it does not erase the remedial question.

G. Third-Party Reliance

The originator may be able to propagate correction without possessing authority to reverse a recipient's independent decision. Responsibilities should be allocated according to authority rather than assumed to reside in one actor.

H. Later-Valid Information

If valid State B appears later and supports continuation, B may justify the consequence prospectively. It does not retroactively become the basis for what happened under invalid State A.

Current justification cannot retroactively validate historical dependency.

I. Participation, Emergency, and Error Contribution

Participation need not be elaborate where it could not affect the result. An affected person may decline participation. Emergency continuation should be narrow, temporary, and reviewable. If the affected person contributed to the original error, that may matter to responsibility or remedy but does not itself answer whether the continuing consequence remains valid.

J. Silent Correction, Vague Revalidation, and Finality

Silent internal correction may be insufficient where the affected person's participation could reveal downstream effects or challenge independent grounds. A vague statement that an outcome remains appropriate may not provide meaningful contestability.

Finality remains legitimate, but closure should distinguish between a consequence that remains because it was substantively revalidated and one that remains because governing law or finality rules bar further reopening.

K. Existing Correction and Automated-Decision Rights

One objection is that existing data-protection tools already do the necessary work. GDPR Article 16 corrects the record. Article 19 communicates that correction to recipients. Article 22, as interpreted in SCHUFA, can reach distributed scoring that functions as an automated decision. On that view, a separate consequence-revalidation framework is unnecessary.

Those tools matter, and this paper relies on them. They still answer a different question. Communicating a corrected record does not determine whether a decision made under the former record remains justified, or whether a continuing restriction may remain. Article 22 constrains certain automated decisions; it does not, by itself, require an institution to revalidate a still-operative consequence after a

later correction. The integration gap remains even if those articles are given their full weight.

L. Revalidation Duties Could Discourage Correction

A second objection is incentive effects. If a material correction obliges an actor to find dependent consequences, reconstruct provenance, and possibly unwind later decisions, the actor may delay or avoid correction.

That risk is real. It is a reason to keep the trigger narrow: a still-operative consequence plus a plausible path of material dependency, not every historical touch. It is also a reason to treat missing provenance as a production problem rather than as automatic invalidity. A framework that punished every correction would not be adopted, and should not be.

M. Explanation Can Collide With Confidentiality

A third objection is that meaningful participation and provenance production can collide with trade secrets, model confidentiality, and third-party data. Dun & Bradstreet Austria makes that collision explicit: the right to meaningful information about the logic involved does not disappear because a trade secret is asserted, but disclosure can be shaped by competing protected interests.

The framework does not require maximum disclosure. It requires enough reconstructability to test continuing authority and enough contestability to challenge that test where participation could affect the result. Confidentiality may limit the form of disclosure. It does not convert non-reconstructability into proof that a consequence remains valid.

N. High-Volume Systems Create Cost and Scale Constraints

A fourth objection is operational scale. Credit scoring, content ranking, fraud flags, and similar high-volume systems generate large numbers of low-visibility consequences. A full revalidation file for every material correction would be unworkable.

Scale is a reason to proportion the process, not to abandon the object. Mechanical recalculation can close some cases without a hearing. Incidental presence can end the inquiry. Serious or continuing consequences, inadequate provenance, and genuinely contestable independent grounds still require more. Cost limits the intensity of the process. It does not erase the question of continuing authority.

O. General Limiting Principle

The framework should be neither maximalist nor cosmetic.

Revalidation should extend as far as material operative dependency extends, and no farther.

X. Institutional Design Implications

Within this proposed architecture:

The framework has direct consequences for how persistent computational systems are built and governed.

Correction cannot remain solely a data-management function where computational state produces durable decisions. Institutions need infrastructure capable of identifying what corrected state affected and determining whether those effects retain authority.

A. Decision Provenance

Conventional provenance asks where information came from. Consequence revalidation also asks what decisions materially relied on it.

Institutions therefore need enough decision provenance to reconstruct the material basis of continuing outcomes. The objective is not permanent preservation of every technical detail. It is sufficient reconstructability to test continuing authority when material state changes.

B. Dependency Discovery and Explicit Revalidation Workflows

A correction interface should not stop at changing a value. Consequential systems should be able to ask where the state is still materially operative.

Institutions should also have explicit workflows for dependency classification, provenance review, participation where useful, revalidation, remedy, propagation, and closure.

C. Authority, Time, and Supersession

Persistent state should not be treated as a context-free fact merely because it remains stored. Systems should preserve enough information about origin, time, authority, supersession, correction, operative status, and pending revalidation.

Superseded state may remain visible without remaining operative.

D. Propagation Tracking and Closure Ownership

Where corrections must travel, institutions should be able to show that propagation occurred and distinguish correction transmitted from consequence resolved.

Consequential correction events also need an identifiable closure authority. Without an owner, correction can become everybody's partial responsibility and nobody's complete responsibility.

E. Continuing Consequences Should Remain Justifiable

An operative consequence should not outlive the institution's ability to identify the basis on which it remains operative.

Where indefinite retention is inappropriate, that may require periodic revalidation, fresh determination, or another domain-appropriate mechanism.

F. Human-Facing Status and Governance Metrics

A person-facing interface should distinguish correction status from consequence status where material. Institutions should also measure not only correction throughput but dependency screening, reconstructability, revalidation time, propagation, remedy, and closure.

G. Vendors, Portability, and Retirement

Contracts and technical interfaces may need to specify correction notification, provenance retention, dependency information, revalidation responsibilities, downstream propagation, supersession handling, and closure evidence.

Portable state should carry enough correction, supersession, and authority information to prevent invalid state from silently regaining force across system boundaries.

Retiring the system does not retire the consequences produced through it.

H. Correction Infrastructure as Agency Infrastructure

Correction systems are not merely compliance infrastructure. Where computational state affects what people can do, correction mechanisms help determine whether people can alter an institutional trajectory after its factual or normative basis changes.

The relevant question is: can a person actually escape the continuing authority of state that the institution itself no longer regards as valid?

XI. Relationship to AI Governance

The consequence-revalidation framework is not a replacement for existing AI governance. It addresses a problem that becomes visible after many familiar governance mechanisms have already done their work.

A system may be monitored, explained, audited, and corrected while a consequence generated under superseded state continues.

A. From System Governance to Consequence Governance

AI governance usually begins with the system: model testing, risk assessment, monitoring, human oversight, and contestability.

Consequence governance asks what the state produced through the system continued to authorize after the system-level issue changed.

That shift matters because consequences can outlive systems.

B. Parallel Lifecycles

The technical lifecycle may resemble Design -> Deployment -> Monitoring -> Modification -> Retirement.

The human consequence lifecycle may resemble State -> Decision -> Consequence -> Persistence -> Revalidation -> Closure.

Those lifecycles intersect but do not terminate together.

C. Output Reversibility Is Not Consequence Recoverability

A flag can be removed while a restriction survives. A score can change while an eligibility decision remains untouched. A model can be corrected while a downstream record continues to encode the earlier conclusion.

Output reversibility is not equivalent to consequence recoverability.

D. Explainability, Contestability, and Human Oversight

Explainability supports dependency analysis; dependency analysis supports revalidation; revalidation determines continuing authority.

Contestability that ends at the source state can succeed formally while failing practically.

Human oversight also needs authority. A reviewer who can see an error but cannot alter the resulting consequence does not complete the governance loop.

E. Audit Error Recovery, Not Only Error Rates

A low error rate says little about what happens to people affected by the errors that do occur.

Mature auditing should ask what happened after the error was discovered: whether dependent decisions were identified, provenance was available, participation was meaningful where needed, consequences were revalidated, remedy occurred, and closure was documented.

This shifts attention from error incidence to error recovery capability.

F. Corrective Continuity as a Governance Property

Corrective continuity is stronger where a valid correction can remain associated with what it supersedes, reach materially dependent decisions, prevent inappropriate future reliance, trigger revalidation, propagate where necessary, and remain legible over time.

The central asymmetry test is: does the architecture preserve the effects of old state more effectively than it preserves valid correction of that state?

G. Persistent Systems Must Govern Remembering and Superseding

Persistence increases the value of memory. It also increases the importance of knowing when memory should cease to govern.

A well-governed persistent system must support both remembering and superseding. Continuity is not maximum retention.

H. Human Agency and Operational Proof

Human agency does not require control over every institutional result. It requires that legitimate change can have practical force.

A person who corrects materially consequential state but remains indefinitely governed by its superseded version has obtained informational correction without consequential recovery.

The operational question is therefore whether an institution can demonstrate that a valid correction can reach the pathways through which superseded state would otherwise continue to matter.

I. Boundary of the Framework

The framework is not a universal liability rule. It does not determine damages, create jurisdiction, displace sector-specific law, eliminate finality, or establish one remedy for every domain.

Its purpose is narrower: to provide an institutional architecture for determining what should happen when consequential computational state changes but effects produced under the former state remain operative.

XII. Conclusion

Persistent computational systems create a temporal governance problem. State about a person can survive the moment in which it was produced, enter later decisions, and continue shaping institutional treatment after the grounds supporting that state have changed.

A correction right is therefore incomplete when the record can change but materially dependent consequences remain untouched.

Existing legal and governance regimes already supply much of the necessary vocabulary: correction, propagation, explanation, reconsideration, participation, restoration, finality, and effective remedy.

What remains fragmented is the connection among those mechanisms once consequential state has acquired continuing institutional force.

This paper has proposed a framework for making that connection.

Correction changes state. Revalidation determines whether effects produced from the former state retain authority.

Correction should trigger revalidation, not automatic reversal.

Failure of reconstructability should not become evidence of continuing validity.

Affected people should have a meaningful opportunity to contest the material basis for continuing consequential effects when participation could matter.

Closure requires resolution of the consequence, not merely correction of the underlying state.

Together, those principles form the lifecycle:

Correction -> Dependency -> Provenance -> Participation -> Revalidation -> Remedy -> Closure.

The framework is deliberately bounded. Not every correction requires reopening. Not every past consequence can be restored. Not every downstream effect remains materially dependent. Independent grounds may support continuation. Finality may constrain review. Third parties may possess separate authority. Time may limit remedy.

The deeper issue is continuity.

Consequential state can persist. Corrections must be capable of persisting as well.

Corrective continuity is the ability of a valid change to remain effective across the pathways through which superseded state would otherwise continue to shape a person's treatment over time.

The governance problem of persistent computational systems is therefore not only whether state may persist, but whether the authority of that state remains revisable when the grounds for its persistence change.

That perspective matters because computational systems increasingly operate through accumulation. Prior classifications shape later judgments. Past decisions become future inputs. Institutional memory becomes part of the environment in which people act.

The ability to remember can produce coordination, efficiency, continuity, and value. But persistence also creates a corresponding governance obligation: the system must be able to recognize when remembered state no longer deserves the authority it once possessed.

Human agency does not require that every adverse consequence disappear when information changes. It requires something more disciplined: valid change must be capable of mattering.

Persistent systems therefore need both memory and revision. They must preserve legitimate continuity while permitting legitimate supersession.

The rule remains bounded. Correction should trigger revalidation of still-operative, materially dependent consequences. Independent grounds, finality, third-party authority, and remedial impossibility may limit what follows.

Those limits should be identified. They should not be supplied by institutional inertia. Meaningful correction cannot end at the record when the authority of the record's consequences continues beyond it.

Notes

1. Regulation (EU) 2024/1689, arts. 20, 72, 2024 O.J. (L 1689) (Artificial Intelligence Act).
2. Regulation (EU) 2016/679, arts. 16, 18-19, 2016 O.J. (L 119) 1 (General Data Protection Regulation).
3. Case C-634/21, SCHUFA Holding (Scoring), ECLI:EU:C:2023:957 (Ct. Just. Dec. 7, 2023).
4. Case C-203/22, Dun & Bradstreet Austria, ECLI:EU:C:2025:117 (Ct. Just. Feb. 27, 2025) (interpreting Regulation (EU) 2016/679, art. 15(1)(h)).
5. *Lee v. Geren*, 480 F. Supp. 2d 198, 206-10 (D.D.C. 2007).
6. 10 U.S.C. § 628(b)(2), (d)(2) (2026).
7. *Roth v. United States*, 378 F.3d 1371 (Fed. Cir. 2004).
8. *Voyageur Outward Bound Sch. v. United States*, 444 F. Supp. 3d 182 (D.D.C. 2020) (agency reconsideration authority and temporal limits; not a person-facing analog).
9. Human Rights Committee, General Comment No. 31, The Nature of the General Legal Obligation Imposed on States Parties to the Covenant, paras. 15-16, U.N. Doc. CCPR/C/21/Rev.1/Add.13 (May 26, 2004).
10. U.N. Office of the High Commissioner for Human Rights, Guiding Principles on Business and Human Rights, princs. 29-31, U.N. Doc. HR/PUB/11/04 (2011).
11. UNESCO, Recommendation on the Ethics of Artificial Intelligence paras. 37-43 (Nov. 23, 2021).