

Human Agency Under Consequential State

Correction, Revalidation, and Recoverable Consequences in Persistent Computational Systems

Micheal Charles Preble

Independent Researcher | ORCID 0009-0005-3906-6028 | September 21, 2026

Study Overview

OD-HAS empirical program

Pilot 1: CFPB - 25 records

Pilot 2: MSPB - 20 chains

Discovery v0.2: search saturation; no qualifying single public corpus identified

Key Findings

- CFPB pilot: 25 records; propagation observable in 25/25; Ground B = 0 YES / 0 NO / 25 UNKNOWN.
- MSPB pilot: 20 chains; Ground B = 1 YES / 19 NO / 0 UNKNOWN; interim authority observable in 0/20; complete clocks in 0/20.
- Discovery v0.2 closed at search saturation: no qualifying single ordinary public corpus was identified under the frozen gates.

Abstract

Persistent computational systems increasingly retain, infer, transmit, and act upon information about people across time. When such information becomes consequential—because it affects access, treatment, authority, resources, opportunities, or participation—the central governance problem is not exhausted by whether the underlying record can be corrected. A correction at the source does not by itself determine what should happen to downstream consequences already produced from the earlier state, whether those consequences may continue under a corrected version of the same ground, whether a materially distinct replacement ground can support them, what authority governs any interim period, or which residual effects remain recoverable after prospective correction. This paper develops Human Agency Under Consequential State (OD-HAS) as a cross-domain framework for analyzing that post-change interval. It distinguishes state change from propagation, propagation from revalidation, same-ground correction from materially distinct replacement grounds, substantive sufficiency from procedural adoption, present authority from historical dependency, and prospective consequence disposition from residual functional effects. The framework defines four state-change types—withdrawal, replacement, annotation, and derived-state change—and separates a corrected same-ground value from a distinct Ground B. Where a materially distinct Ground B is substantively sufficient but not procedurally adopted, the framework identifies a narrow REBASE-PENDING condition rather than treating substantive readiness as self-executing authority. The Rule A and Rule B sections in this paper are a conceptual exposition of the governing distinctions, not the executable coding specification. Case-level empirical classifications were produced under OD-HAS-01-SUT-001 and its frozen supplements; that instrument is not reproduced in this manuscript. The paper also reports the first empirical contact of the framework with real public records. Pilot 1 used 25 Consumer Financial Protection Bureau complaint records selected for observable state change and post-change consequences. Those records were rich in propagation and continuing effects but uniformly unable to resolve Ground B: 0 YES / 0 NO / 25 UNKNOWN. Pilot 2 used 20 Merit Systems Protection Board procedural chains selected for basis change, reassessment, and later disposition. In that corpus, Ground B became fully decidable: 1 YES / 19 NO / 0 UNKNOWN, while interim authority remained unobservable in 20/20 chains and resolution clocks were complete in 0/20. Across both pilots, no new Rule A category, Rule B category, recurrent O4 misfit, mid-pilot patch, or SUT change was required. A subsequent frozen discovery program searched for a single ordinary public corpus jointly exposing changed-basis reassessment, a live interim consequence, routinely public causal reasons, and a consequence-linked authority deadline with an observable lapse effect. No qualifying corpus was identified under the precommitted Discovery v0.2 gates. Across the record systems examined, public observability of basis revision repeatedly separated from public observability of interim authority and binding timing. This is a bounded finding about public-record structure, not institutional reality and not a claim that no qualifying corpus exists. The resulting contribution is therefore both conceptual and methodological: OD-HAS specifies a governable post-change sequence, while the empirical work shows that testing the sequence depends on record ecologies capable of exposing different parts of it.

Keywords: consequential state; human agency; AI governance; correction; revalidation; public records; algorithmic recourse; provenance; persistent computational systems

1. Introduction

Computational systems increasingly carry consequential information forward. A person may be represented by a credit tradeline, screening result, fraud flag, eligibility status, risk classification, professional finding, employment charge, administrative determination, identity relation, or derived inference. This problem sits within a broader human-centered and sociotechnical literature that treats human control, accountability, and institutional context as part of system design rather than as external add-ons (Selbst et al., 2019; Shneiderman, 2020; Tabassi, 2023). That state can influence what happens next. It can alter whether a person receives credit, keeps a job, passes a screening process, remains subject to a restriction, gains access to a service, or must overcome an institutional burden created by an earlier classification.

Much of the existing correction problem is framed around the accuracy of the state itself. That framing is necessary but incomplete. A source record can be corrected while a downstream consequence remains in place. A charge can be withdrawn while a penalty survives on remaining grounds. A report can be deleted by one actor while another actor continues to display or rely on it. A mistaken predicate can be repaired while a derived classification remains unchanged. A temporary consequence can remain operative during reconsideration even after part of its original basis has changed. In each of these situations, the important question is not only whether state changed, but what authority remains for the consequence after the change.

OD-HAS addresses this interval. Its object is the period after consequential state materially changes but before the resulting consequences have necessarily been retired, modified, rebased, or otherwise resolved. The framework asks which actors control the continuing consequence, what ground originally supported it, whether that ground still survives, whether a corrected version of the same ground remains sufficient, whether a materially distinct Ground B exists, whether that Ground B has been adopted through the procedure required to make it operative, what authority governs any interim continuation, and which functional effects remain after prospective correction. The framework is intentionally cross-domain. It is not offered as a substitute for sector-specific law, nor does it attempt to collapse distinct legal doctrines into one rule. Privacy law, consumer reporting law, administrative law, licensing systems, employment procedures, sanctions regimes, records law, and sector-specific correction mechanisms all provide partial answers to different portions of the problem. OD-HAS instead supplies a common analytic structure for asking what must be known before a live consequence can be treated as currently warranted after its basis has changed.

The empirical work reported here changes the status of the paper. The framework is no longer purely pre-empirical. It has now been applied, under a frozen specification and explicit missingness rules, to two heterogeneous public-record corpora. The result is neither universal validation nor framework failure. Instead, the pilots show that different public record environments expose different portions of the post-change chain. Complaint narratives can expose correction, propagation, and continuing effects while concealing internal institutional grounds. Procedure-rich adjudicatory records can expose grounds, basis changes, reassessment, and later disposition while still omitting the authority and timing governing an interim state.

A subsequent discovery program attempted to identify a third ordinary public corpus that would expose the missing conjunction. That search closed at saturation without identifying a qualifying single corpus. The repeated separation between basis-revision visibility and authority-and-clock visibility is therefore itself part of the empirical result. OD-HAS now makes two claims of different kinds. First, it proposes a cross-domain structure for reasoning about continuing consequences after consequential state changes. Second, its empirical program shows that the observability of that structure is strongly conditioned by the architecture of the public record through which it is studied.

2. Consequential State

2.1 Definition

Consequential State (CS) is machine-actionable or institutionally actionable information, classification, inference, status, score, annotation, credential, or representation about a person that can materially affect access, treatment, authority, resources, opportunities, or participation. The definition is deliberately functional. A piece of information is not consequential merely because it exists, is stored, or is personally meaningful. It becomes consequential for this framework when an actor can use it, directly or through a system, to produce or maintain a materially relevant effect. That actor may be the original collector, a downstream institution, a screening service, an adjudicator, an employer, a financial intermediary, a regulator, or another consequence controller.

The state may be explicit or derived. It may originate in a source record, a third-party report, an institutional finding, a rule application, or a downstream inference. The framework does not assume that all consequential state is inaccurate, harmful, or illegitimate. Its concern begins when a material change occurs and the continued status of an already-existing consequence must be assessed.

2.2 State is not consequence

A central distinction is:

state ≠ consequence.

A record can change without the consequence changing. A consequence can change even when the underlying state remains materially the same. This matters because correction systems often focus on the state-bearing actor while the operative consequence is controlled elsewhere. If Actor A corrects a source record and Actor B still controls a live consequence, the fact of correction does not itself answer whether Actor B must retire, narrow, reassess, or continue that consequence. The framework therefore treats source-state correction and consequence disposition as separate analytical events.

2.3 Historical dependency and current warrant

A second distinction separates historical explanation from present justification. Let:

D_{t_0}

denote the basis on which a consequence was originally imposed or maintained at the original decision time, and let:

W_{t_1}

denote the warrant available at the later reassessment time after material state change. A consequence may have been historically dependent on a state that no longer exists in the same form. The fact that the original decision was validly made does not, by itself, establish that the consequence remains currently warranted. Conversely, the fact that one part of an original ground changed does not automatically establish that no remaining current warrant exists.

OD-HAS therefore evaluates original dependency and current warrant separately.

3. The Post-Change Lifecycle

The framework organizes the relevant sequence as:



These stages are related but non-equivalent.

3.1 State change

The first question is what changed. A dispute, challenge, or allegation of inaccuracy is not itself a state-change event. The framework requires an observable material change such as deletion, withdrawal, correction, replacement, removal, reinsertion, changed status, acknowledged

correction, changed derived state, or another alteration with possible consequence relevance.

3.2 Propagation

The second question is whether the change reaches actors or systems that hold or rely on related state. Propagation concerns information movement and update. It may involve source-to-recipient correction, cross-actor notice, record synchronization, downstream deletion, recomputation, or other mechanisms. Propagation is necessary in many systems but not sufficient. A recipient may learn that a source state changed while still needing to decide whether its own consequence should continue.

3.3 Revalidation

Revalidation asks whether a consequence-controlling actor must reassess the live consequence in light of the changed state. The question is not merely whether the actor has notice, but whether the present warrant for the consequence remains sufficient under the actor's governing rule, purpose, authority, and procedure. This is the central move from information correction to consequence governance.

3.4 Disposition

After revalidation, the consequence may be:

- RETIRE — no current sufficient basis remains for the consequence.
- MODIFY — a remaining same-ground basis supports a narrower or otherwise altered consequence.
- REBASE — a materially distinct replacement Ground B has been properly adopted and supports the consequence.
- UNRESOLVED — the available evidence is insufficient to determine a valid disposition.

REBASE-PENDING is not itself a final disposition. It is a procedural state defined below.

3.5 Recoverable effects

Prospective correction does not necessarily erase what has already happened. A corrected report may leave a lost opportunity. A withdrawn charge may leave delay, cost, stigma, or institutional burden. A screening correction may arrive after a decision window closed. A changed administrative finding may not restore time, access, or trust.

OD-HAS therefore treats recoverable functional effects as a separate dependent variable. The existence and recoverability of those effects must be studied rather than assumed. They are not foundational primitives of the framework and they are not automatically equivalent to legal damages.

4. Rule A: Corrected-State Treatment

Rule A governs the treatment of changed state before any materially distinct replacement ground is considered.

4.1 State-change types

The frozen state-change vocabulary contains four primary types:

- **WITHDRAWAL** — the prior state is removed, reversed, vacated, deleted, or otherwise no longer asserted in its former operative form.
- **REPLACEMENT** — the prior state is replaced by a corrected or revised value concerning the same underlying ground.
- **ANNOTATION** — a note, dispute marker, explanation, caveat, or contextual statement is added while the underlying state remains otherwise operative.
- **DERIVED-STATE CHANGE** — a downstream inference, score, classification, or derived representation changes.

Observed cases may contain more than one of these operations. **MULTIPLE** is therefore a combination code rather than a new substantive state-change type.

4.2 Withdrawal

A withdrawn state does not remain present authority merely because it once existed. Historical existence can remain relevant to audit, provenance, chronology, or other lawful purposes, but history is not self-validating present warrant.

4.3 Corrected same-ground state

Where a value changes from:

$$v \rightarrow v',$$

the corrected value (v') remains part of the same underlying ground unless the later basis is materially distinct in kind rather than merely corrected in value. This distinction is critical. A corrected amount, corrected date, narrowed specification, surviving charge, or revised finding may continue to belong to the same ground. OD-HAS does not convert every correction into a new Ground B.

4.4 Same-ground warrant

A corrected same-ground value is relevant to current consequence only where an applicable governing rule is in force at reassessment, applies to the consequence and authorized purpose, the actor has authority to apply it, and the corrected state satisfies the rule. A generally applicable policy in force at reassessment may establish a same-ground floor. A file-specific memorandum, after-the-fact explanation, or unadopted alternative rationale does not automatically become either that floor or a materially distinct Ground B.

4.5 Annotation

An annotation is not independently sufficient warrant unless an applicable rule gives that annotation operative effect. A dispute marker may change how a record should be interpreted without becoming a replacement basis for the underlying consequence.

4.6 Derived state

When a source state changes, a downstream inference or classification does not automatically become corrected. The framework separately asks about lineage, recomputation, current inputs, current warrant, and continuing use. This separates:

source correction

from:

derived-state correction.

4.7 Base disposition

Rule A determines the base disposition before distinct Ground B analysis. If corrected same-ground state supports a narrower consequence, the base disposition may be **MODIFY**. If no current same-ground warrant remains, the base disposition may be **RETIRE**. A later Ground B analysis cannot erase this ordering.

5. Rule B: Distinct Ground B and REBASE-PENDING

Rule B addresses a narrower circumstance: the possibility that the original or corrected same-ground basis no longer supports the full consequence, but a materially distinct replacement basis exists.

5.1 Distinctness

A candidate Ground B must be materially distinct from the corrected version of Ground A. Corrected (v') is never, by itself, Ground B. The distinction is functional rather than merely verbal. Rewording the original ground, adding detail to the same underlying allegation, or relying on a surviving portion of the same charge structure does not necessarily create Ground B.

5.2 REBASE-PENDING trigger

A case enters **REBASE-PENDING** only where all four predicates are satisfied:

1. a materially distinct Ground B exists;
2. Ground B is substantively sufficient;
3. Ground B has not been procedurally adopted;
4. Ground B would support substantially the same consequence or increment.

Where Ground B would support a more severe consequence, the additional severity is a new action rather than a continuation through REBASE-PENDING.

5.3 Substantive sufficiency is not adoption

A central principle is:

substantive sufficiency $\neq$ procedural adoption.

A replacement ground can be factually or legally capable of supporting a consequence while remaining unavailable as present authority because the actor has not completed the procedure required to make that ground operative. This prevents an institution from preserving a consequence by silently switching rationales after the fact.

5.4 Interim floor

Where a procedurally operative same-ground floor remains, interim continuation may be narrowed to that floor. The existence of a substantively sufficient but unadopted Ground B does not itself authorize continuation above the already-operative floor. Where no such floor exists, the unsupported consequence or unsupported increment should not be preserved merely because replacement reasoning is available.

5.5 Interim safeguards

A temporary targeted safeguard is analytically distinct from preserving the old consequence. Such a safeguard must already be authorized, risk-targeted, narrower than the unsupported consequence, temporary, documented, and not functionally recreate the very consequence whose warrant has failed.

5.6 Exceptions and authority

Certain systems may contain specific authority that preserves an interim state during reconsideration, such as emergency authority, a court order maintaining status quo, or a statute or regulation expressly continuing the consequence pending review. The relevant question is whether such authority actually governs the case.

5.7 Resolution clock

REBASE-PENDING cannot remain indefinitely open without a governing timing rule. A conforming record should identify the adoption procedure, responsible actor, operative interim state, and a calendared deadline or binding authority that supplies timing. The empirical work in this paper shows that this final layer—interim authority plus resolution clock—is particularly difficult to observe in ordinary public records.

6. Warrant as a Relational Constraint

OD-HAS treats warrant as relational:

$$W = W(S, C, P, A, t),$$

where: (S) is the currently relevant state; (C) is the consequence; (P) is the authorized purpose; (A) is the actor with relevant authority; (t) is the reassessment time. The expression is not proposed as a scalar score or a complete quantitative theory. It is a constraint structure. A state that could justify one consequence for one purpose under one actor's authority may not justify a different consequence, purpose, actor, or time.

This formulation blocks several shortcuts. Historical truth does not establish present authorization. Factual sufficiency does not establish procedural adoption. Availability of a rationale does not establish authority to use it. A record's persistence does not make its continuing use legitimate merely because the record is technically accessible.

7. Human Agency and Recoverability

The paper's title emphasizes human agency because consequential computational state can constrain the practical range of actions available to a person. A status, classification, flag, score, or administrative finding can affect not only what an institution knows, but what a person can realistically do next. Correction therefore has at least two levels. The first is representational correction: the relevant state is repaired, withdrawn, annotated, or recomputed.

The second is functional correction: the consequences produced from the former state are reconsidered, and where appropriate retired, narrowed, rebased, or otherwise resolved. The gap between these levels is where residual effects arise. Those effects may include lost access, lost time, denied opportunities, procedural burden, reputational consequences, duplicated effort, repeated proof obligations, or institutional inertia. OD-HAS does not assume that every residual effect is remediable or attributable to one actor. It identifies the recovery question as analytically separate from prospective correction.

Human agency is therefore not defined here as unconstrained choice. It concerns whether a person can meaningfully escape, correct, contest, or recover from consequences that persist after their stated basis has materially changed. This concern is adjacent to algorithmic-recourse work, which asks whether a person can take meaningful action to alter a consequential model outcome, while OD-HAS focuses on what happens when the operative institutional state itself has already changed (Ustun et al., 2019).

8. Prior-Art Reduction and Residual Contribution

The underlying problem intersects with existing legal and technical mechanisms. Data-protection systems contain rights of rectification, restriction, and recipient notification, including the GDPR's rectification, restriction, and downstream-notification provisions (Regulation (EU) 2016/679, arts. 16, 18-19). Consumer-reporting systems contain dispute, reinvestigation, correction, and reinsertion mechanisms, including the Fair Credit Reporting Act's reinvestigation and deletion-or-modification structure (15 U.S.C. § 1681i). Administrative law contains reconsideration, vacatur, remand, reopening, and distinctions between procedural and substantive authority. Sanctions and licensing systems can require reassessment when predicates change. Records systems distinguish historical preservation from current operative use. Technical unlearning and model-editing work likewise shows that changing a source datum and changing a learned or derived representation are distinct technical operations rather than automatic equivalents (Meng et al., 2022; Shaik et al., 2025).

OD-HAS does not claim those mechanisms as new. Its residual contribution is the cross-domain separation of five questions that are frequently collapsed:

- What state changed?
- Where did that change propagate?
- Which actor still controls a live consequence, and what current warrant remains?
- If a distinct replacement ground exists, is it merely sufficient or actually adopted?
- What prospective and residual effects remain after the consequence is reassessed?

The framework's distinctive middle interval is therefore the period after consequential state changes but before authority over the continuing consequence has been fully resolved.

9. From Conceptual Framework to Empirical Contact

9.1 Research objective

The empirical program did not attempt to estimate prevalence, prove legal compliance, or validate OD-HAS universally. Its objective was narrower: determine whether the framework's constructs could be observed in real public records under a frozen specification without rewriting categories around individual cases. The empirical program separated four layers: acquisition reproducibility — whether the public source universe and retained records could be reconstructed; screening reproducibility — whether inclusion criteria could be applied without SUT classification; coding reproducibility — whether the frozen

instrument and SUT could classify the included records; construct observability — whether the public record actually exposed the facts required by the construct. Missingness was treated as a result. UNKNOWN was not converted into NO, and public silence was not treated as proof of institutional absence.

9.2 Frozen specification

Across Pilots 1 and 2: no Rule A category was added; no Rule B category was added; no O4 recurrent misfit occurred; no mid-pilot patch was applied; no SUT change was made. This matters because empirical contact was allowed to reveal record limitations rather than forcing the framework to produce determinate outputs.

10. Pilot 1 — CFPB Complaint Narratives

10.1 Corpus rationale

Pilot 1 used Consumer Financial Protection Bureau complaint narratives because those records can expose source disputes, corrections, deletions, reporting differences, reinsertion, continuing downstream listings, and consumer-described residual effects. The CFPB's public materials define Complaint ID as the unique identifier for a complaint and preserve previously published narratives in a 2011-2026 archive, while also emphasizing that complaint narratives are unverified, one-sided allegations rather than adjudicated findings (Consumer Financial Protection Bureau [CFPB], 2026a, 2026b, n.d.). The acquisition process produced a large metadata-first candidate pool and then applied frozen product, issue, narrative, keyword, chronology, and screening rules. The final included corpus contained 25 complaint records with an identifiable consequential state, live consequence, material state-change event, post-change information, and sufficient ordinal chronology.

The inclusion rule intentionally selected post-change problems. It therefore cannot support prevalence claims about the broader complaint database, consumer-reporting market, successful correction rate, or general incidence of residual effects.

10.2 Pilot 1 result

All 25 included cases completed coding as partially codable records:

O1 — Fully codable	0
O2 — Partially codable	25
O3 — Public-data weak	0
O4 — Misfit	0
Propagation was observable in:	

25 / 25 cases. Revalidation was observable in: 9 / 25 cases. Ground B was:

YES 0
 NO 0
 UNKNOWN 25
 REBASE-PENDING was:

YES 0
 NO 0
 UNKNOWN 25
 The Rule A state-change distribution was:

WITHDRAWAL 1
 REPLACEMENT 2
 MULTIPLE 22
 ANNOTATION 0
 DERIVED 0
 UNOBSERVABLE 0
 Residual effects were observable across the included post-change problem corpus, but that result is conditioned by the inclusion criteria and cannot be generalized as a market rate.

10.3 What CFPB exposed

The complaint narratives were strong for: material state change; cross-actor propagation; continued reporting or related consequences; post-change persistence; residual functional effects. They were weak for: internal institutional grounds; governing warrant text; distinct Ground B; procedural adoption; interim authority; resolution clocks. The appropriate Pilot 1 decision was therefore HOLD.

10.4 Interpretation

The key finding from Pilot 1 is not that Rule B failed. It is that the public complaint layer could not distinguish the absence of Ground B from nonobservability of Ground B. Formally:

UNKNOWN ≠ NO.

That result justified moving to a deliberately procedure-rich second corpus rather than changing the framework.

11. Pilot 2 — MSPB Procedural Chains

11.1 Corpus rationale

Pilot 2 used public Merit Systems Protection Board adverse-action and remand records because those decisions can expose charges, failed or surviving specifications, penalty reasoning, formal remand instructions, responsible actors, procedural reassessment, and later disposition. MSPB publicly maintains separate precedential and nonprecedential decision collections, including five-year listings and advanced search functions that support docket-based reconstruction of Board-level chains (U.S. Merit Systems Protection Board [MSPB], n.d.-a, n.d.-b). The acquisition design separated Board-recited procedural completeness from availability of underlying lower-stage documents. A candidate frame was

constructed from public § 7512/§ 752 adverse-action records, using a primary B/M frame and a precommitted I-remand extension when the primary frame produced fewer than the target 20 included chains.

The final corpus contained:

17 primary B/M-frame chains
 3 I-remand extension chains
 20 total
 The extension provenance remains distinct in reporting.

11.2 Pilot 2 result

All 20 chains again coded as partially codable:

O1 0
 O2 20
 O3 0
 O4 0
 Ground B was fully decidable:

YES 1
 NO 19
 UNKNOWN 0
 Ground B substantive sufficiency was relevant in one chain and coded YES.

Procedural adoption for that Ground B was:

YES 0
 NO 1
 UNKNOWN 0
 N/A 19
 REBASE-PENDING was:

YES 0
 NO 1
 UNKNOWN 0
 N/A 19
 Interim authority was:

OBSERVABLE 0
 UNKNOWN 20
 Resolution-clock observability was:

COMPLETE 0
 PARTIAL 1
 NONE OBSERVED 19
 Final dispositions were reported as:

RETIRE 3
 MODIFY 9
 REBASE 0
 UNRESOLVED 7
 OTHER/UNKNOWN 1

11.3 Ground B observability

The contrast with Pilot 1 is the strongest empirical result in the two-pilot phase. CFPB: Ground B 0 YES / 0 NO / 25

UNKNOWN MSPB: Ground B 1 YES / 19 NO / 0
UNKNOWN The construct therefore did not remain intrinsically unobservable. Its observability changed with the public record environment. A bounded formulation is:

Ground B observability

=

f(record structure, procedural recitation, available institutional facts).

This does not establish how common Ground B is in MSPB proceedings, administrative law, employment disputes, or institutional decision-making generally.

11.4 Same-ground persistence

In the selected MSPB sample, Ground B = NO in 19/20 chains. Within this corpus, basis change usually appeared as surviving charges, surviving specifications, narrowed findings, or penalty reassessment within the original ground structure rather than as a materially distinct replacement ground. That observation is sample-specific. It is not a prevalence estimate.

11.5 Diagnostic Ground B-positive case

One MSPB chain exposed a materially distinct candidate Ground B, allowed substantive sufficiency to be assessed, and showed that the distinct ground had not been procedurally adopted at the relevant point. Under the frozen SUT, REBASE-PENDING was coded NO and the disposition was RETIRE. The manuscript-only record is insufficient to determine why that classification was NO. At least two possibilities remain open from the published text: the fourth conceptual predicate in §5.2 may not have been satisfied in the diagnostic chain, or the frozen SUT may contain a further controlling decision rule not reproduced in this manuscript. This paper does not choose between those possibilities.

That chain is retained as a diagnostic example only. It does not represent MSPB practice generally and is not used as a showcase for the framework. The reported classification should therefore be read as an empirical SUT output, not as a worked derivation from §§5.1–5.7.

11.6 Pilot 2 HOLD

Pilot 2 closed at HOLD because the corpus solved one missingness problem while preserving another. Grounds and basis changes became decidable. Interim authority and binding timing did not. The most precise summary is: Procedure-rich is not clock-rich.

12. Cross-Pilot Findings

12.1 Complementary observability

The two corpora exposed different layers of the same frozen specification.

Construct layer	CFPB	MSPB
Material state change	strong	strong
Cross-actor propagation	strong	limited / case-dependent
Continuing consequence	strong	strong
Revalidation procedure	limited	strong
Original ground	weak	strong
Same-ground analysis	weak	strong
Ground B	opaque	decidable
Procedural adoption	opaque	partly decidable
Interim authority	opaque	opaque
Resolution clock	opaque	opaque
Later disposition	complaint-conditioned	strong
Residual functional effect	strong	limited
The result is not that one corpus was better. Each record ecology exposed a different portion of the post-change chain.		

12.2 Framework stability under contact

Across both pilots:

New Rule A categories	0
New Rule B categories	0
O4 misfits	0
Mid-pilot patches	0
SUT changes	0
This does not validate OD-HAS universally. It supports a narrower proposition: the frozen specification could contact two heterogeneous real public-record environments without recurrent category collapse, while its missingness rules revealed different evidence boundaries in each corpus.	

12.3 Public-record structure as measurement condition

The pilots establish that construct observability cannot be treated as independent of the record environment. A complaint narrative, Board decision, interim order, final disposition, internal memorandum, and docket record each preserve different institutional facts. A missing construct can therefore have at least three meanings: the construct is absent; the construct exists but is not present in the particular case; the relevant facts are not exposed by the public record. OD-HAS does not permit the third condition to be silently recoded as the first.

13. OD-HAS-03 Discovery v0.2: Negative Discovery Result

13.1 Why a third corpus was sought

After Pilot 2, the unresolved empirical frontier was specific:

interim authority

+

binding resolution timing.

A third corpus was not intended merely to add more cases. It was intended to expose the combination that Pilots 1 and 2 could not jointly observe. The search target eventually hardened into four pre-probe requirements: a materially changed predicate must trigger mandatory or formally required reassessment; the consequence must remain live during reassessment; ordinary public decisions must disclose the causal reasons for the interim outcome; a consequence-linked authority deadline must have an observable lapse, nonrenewal, nonaction, or missed-deadline effect. Candidates clearing all four gates then had to satisfy a five-chain density test using ordinary, non-landmark, non-positive-seeded cases.

13.2 Discovery failures were typed, not collapsed

The search examined multiple institutional families and rejected candidates for different reasons. Some systems were authority-and-clock rich but basis-change poor. Others required changed-basis review but kept the reasons private. Others published rich reasons and hard deadlines but made new evidence discretionary rather than the trigger. Some had event-based endpoints rather than true deadline-linked lapse. One system, NMC, appeared to supply all four structural elements but still failed the ordinary-case density threshold.

These were not treated as interchangeable failures.

13.3 SEC and BIS

SEC Rule 102(e)(3) temporary suspensions passed an authority-and-clock discovery test but produced only 1/3 ordinary chains with basis-change compatibility. BIS § 766.24 Temporary Denial Orders appeared stronger in a three-chain discovery probe and therefore advanced to an exhaustive Federal Register acquisition. The final acquisition contained:

136 retained Federal Register documents
104 finalized procedural chains
Only:

1 / 104 chains exposed the full factual conjunction of material predicate change, live TDO, authority, timing, and expiration effect. This was a critical methodological correction. Discovery feasibility did not equal corpus-level density.

13.4 TSBDE and NMC

Texas dental-board emergency suspensions exposed authority, timing, and automatic reinstatement architecture in 5/5 tested chains, but live material basis change was 0/5 (Texas State Board of Dental Examiners [TSBDE], 2026).

NMC interim-order reviews supplied the strongest combined architecture found during directed discovery: new evidence could trigger review, the consequence remained live, reasoned public decisions were available, and the interim order had defined duration/extension mechanics. NMC guidance states that new relevant evidence can require an early review and that review panels reconsider whether an interim restriction remains necessary and proportionate in light of changed circumstances (Nursing and Midwifery Council [NMC], 2025). Yet an ordinary chronological January 2026 seed could produce at most 1/5 potential live material basis-change chains, below the frozen 3/5 threshold.

Exceptional litigation could not rescue that density result.

13.5 Saturation

The search ultimately closed at SEARCH SATURATION. The frozen result is: No qualifying single ordinary public corpus was identified under OD-HAS-03 Discovery v0.2. That statement is intentionally bounded. It does not mean that no such corpus exists. It does not mean institutions lack the relevant mechanisms internally. It does not mean Rule B is unobservable in principle. It means the directed discovery program, under its frozen gates, did not identify a public corpus with sufficient ordinary-case density to support the planned single-corpus Pilot 3.

13.6 Recurring observability partition

Across the examined record environments, a recurring partition appeared: CFPB exposed state effects and propagation but not internal grounds or clocks. MSPB exposed grounds, basis change, reassessment, and disposition but not interim authority or clocks. SEC, BIS, and TSBDE exposed authority and timing but insufficient live basis revision. EWC and HCPTS exposed changed-basis review architecture and authority boundaries but not routinely public causal reasons (Education Workforce Council [EWC], 2026; Health and Care Professions Tribunal Service [HCPTS], n.d.; Health Professions Order 2001, art. 31).

GPhC, FAA/NTSB, and bankruptcy § 362(c)(3) exposed public reasons and hard timing but not changed predicate as the mandatory reassessment trigger. NMC exposed all four structural elements but failed the ordinary density threshold. The resulting bounded empirical proposition is: In the public institutional record systems examined under OD-HAS-03 Discovery v0.2, observability of basis revision and observability of interim authority and binding timing repeatedly separated rather than co-occurring at sufficient ordinary-case density for the planned single-corpus Pilot 3.

14. Discussion

14.1 The empirical result is about observability, not hidden institutional reality

The most important discipline in interpreting these results is to preserve the difference between record structure and institutional process. A public complaint may omit an internal rationale that exists. A Board decision may summarize a remand without reproducing the interim authority that governed implementation. A public suspension page may state an outcome while withholding the reasons. A regulation may supply an authority clock while ordinary case files rarely contain a material changed basis.

Therefore:

not publicly observed
≠
institutionally absent.

The empirical program is best understood as mapping which parts of the framework can be tested in which public record environments.

14.2 Correction and consequence governance are distinct measurement problems

Pilot 1 strongly exposed correction and propagation. Pilot 2 exposed basis and reassessment. The third-corpus search showed that authority and clock observability often appears in yet another record family. This reinforces the conceptual premise that correction is not one event. A system may be transparent about the source-state correction while opaque about who controls the downstream consequence. Another may explain the grounds and later disposition while omitting the interim authority. A third may publish timing rules while providing little evidence of actual basis revision.

The empirical separation therefore mirrors the analytic separation in the framework.

14.3 Ground B is not a generic alternative rationale

The MSPB result is useful because it demonstrates how rarely a distinct Ground B appeared once the same-ground boundary was enforced. In 19/20 selected chains, the remaining structure was coded within the original ground rather than as a materially distinct replacement basis. This is not evidence that Ground B is rare in all institutional systems. It is evidence that the distinction between corrected or surviving same-ground state and a distinct replacement ground is operationally important. Without that distinction, ordinary narrowing or penalty reassessment could be mislabeled as rebasing.

14.4 Rule B's empirical bottleneck is not conceptual alone

The public-record search suggests that the most difficult portion of Rule B to test is the conjunction of: material basis change; live consequence; replacement-ground analysis; interim authority; binding resolution timing; public causal reasoning. Different institutional systems often make some of those facts public while keeping others private, implicit, distributed, or rare. That observation should temper claims about validation. The absence of a single public corpus capable of testing the whole sequence means some Rule B predicates remain only partially empirically contacted.

14.5 Negative discovery as result

The search-saturation finding should not be described as an unsuccessful literature or corpus search. The gates were intentionally strengthened after early discovery examples proved too permissive. BIS was particularly important: a promising 2/3 probe became only 1/104 full-conjunction chains under exhaustive acquisition. That experience justified a density gate and ultimately supported stopping. Continuing to search regulators after the same failure pattern repeated would have increased the risk of positive-example hunting.

The negative discovery result is therefore methodological evidence about the relationship between a construct and its public measurement environment.

15. Limitations

The empirical corpora are small, selected, and designed for construct observability rather than representativeness. No prevalence estimate follows from the 25 CFPB records or 20 MSPB chains. The CFPB sample was explicitly conditioned on observable state change and post-change information, so its residual-effect pattern cannot be generalized to all complaints or to the consumer-reporting market. The MSPB sample was conditioned on public procedural richness and later supplemented by a precommitted I-remand extension to reach the target corpus size.

Complaint narratives are allegations and should not be treated as adjudicated truth. MSPB Board decisions may summarize lower-stage proceedings without publishing every underlying document. Public administrative records can omit internal memoranda, implementation instructions, hearing materials, source letters, timing orders, or other facts that exist within the institution. The discovery program was broad but not exhaustive. It examined multiple jurisdictions and institutional families under a progressively frozen gate, but it did not prove universal nonexistence of a qualifying corpus. Its saturation decision is a stopping rule for this research branch, not a claim about all public institutions.

The manuscript does not reproduce OD-HAS-01-SUT-001 or its frozen supplements. Accordingly, the empirical coding tables are source-controlled outputs of that instrument, but not every case-level classification can be independently regenerated from the narrative Rule A and Rule B exposition printed here. This is a manuscript-level reproducibility limitation, not evidence that the frozen SUT or empirical cells have changed. The SUT remained unchanged through the pilots and discovery program. That stability is useful, but it should not be overstated. The framework has not been independently reconstructed across all implementation details, has not been validated as law, has not been shown to predict institutional outcomes, and has not yet been tested against a single ordinary public corpus exposing the full Rule B authority-and-clock chain.

Recoverable functional effects also remain underdeveloped empirically. Pilot 1 made them visible because the corpus was selected for post-change problems. A future study would need a different sampling design to estimate their frequency, severity, duration, or recoverability.

16. Implications for System and Institutional Design

The empirical findings reinforce several design implications without converting them into universal legal duties. First, systems that permit consequential state to persist should preserve enough provenance to identify which consequences depend on which grounds. Without lineage, a correction cannot reliably trigger downstream reassessment. Second, correction infrastructure should distinguish source repair from derived-state repair. A corrected source value should not be assumed to recompute every downstream classification, score, or operational decision.

Third, consequence-controlling actors should maintain explicit records of current warrant. If an institution can continue a consequence after material state change, the relevant ground, authority, purpose, and procedure should be inspectable enough to determine whether that continuation is same-ground, properly rebased, temporarily authorized, or unsupported. Fourth, interim authority should not be implicit. Where a consequence remains live during reassessment, the governing authority and timing should be explicit. The empirical difficulty of observing these facts in public records is itself a warning: institutions may need more legible interim-state records if meaningful audit is expected.

Fifth, correction systems should distinguish prospective correction from functional recovery. Updating a record may stop future reliance while leaving past consequences unresolved. These design implications follow from the architecture and the observed measurement gaps. They

remain proposals for governance, not claims that existing law universally imposes them.

17. Research State and Deferred Fork

The empirical branch now closes in the following state:

Pilot 1 — CFPB
CLOSED — HOLD

Pilot 2 — MSPB
CLOSED — HOLD

Pilot 3
NO QUALIFYING SINGLE PUBLIC CORPUS IDENTIFIED

- OD-HAS-03 Discovery v0.2
FROZEN — SEARCH SATURATION

Acquisition NOT AUTHORIZED Screening NOT AUTHORIZED Coding NOT AUTHORIZED SUT UNCHANGED Rule A UNCHANGED Rule B UNCHANGED Architecture CLOSED One consequential research question remains deliberately deferred: Should a future empirical program retain the requirement that one corpus expose the entire conjunction, or should a new design test complementary OD-HAS layers across different record ecologies? The first path leaves Pilot 3 dormant until an independently identified qualifying corpus appears. The second path would be a new research design, not a repair or continuation of Pilot 3. It would change the unit of evidence and require a new identifier, protocol, falsifiers, and freeze process.

No such redesign is made in this paper.

18. Conclusion

Consequential state is not self-retiring. When information that helped produce a live consequence is corrected, withdrawn, narrowed, replaced, or otherwise materially changed, the consequence does not answer for itself. Someone or some institution still controls whether it persists, under what ground, for what purpose, through what authority, and for how long. OD-HAS proposes a structured way to analyze that interval. Rule A separates withdrawal, replacement, annotation, and derived-state change and asks what same-ground warrant remains after correction. Rule B separates a materially distinct Ground B from corrected same-ground state, distinguishes substantive sufficiency from procedural adoption, and identifies the narrow REBASE-PENDING interval in which a distinct replacement ground may be sufficient but not yet operative. The framework then separates final disposition from residual functional effects.

The empirical program provides a more disciplined basis for interpreting those constructs. CFPB complaint narratives exposed state change, propagation, continuing reported consequences, and residual effects, but not the internal procedural predicates required to resolve Ground B. MSPB procedural chains exposed original grounds, basis changes, reassessment, and later disposition and made Ground B decidable in all 20 selected chains, but still did not expose the interim authority and binding timing needed to complete the Rule B chain. The subsequent discovery program did not find a single ordinary public corpus that jointly exposed changed-basis reassessment, a live interim consequence, routinely public causal reasons, and a consequence-linked authority deadline at sufficient ordinary-case density. Rather than weakening the criteria until a corpus passed, the program closed at saturation.

The resulting finding is narrow but important. In the public institutional record systems examined, basis-revision visibility and authority-and-clock visibility repeatedly appeared in different places. That does not show that institutions lack an integrated internal process. It shows that public accountability for the post-change interval is fragmented across record ecologies.

The paper therefore ends without claiming universal validation and without repairing the architecture around missing public evidence. The framework remains frozen. Pilots 1 and 2 remain HOLD. Pilot 3 does not begin. The next substantive research decision—whether to retain the single-corpus requirement or design a new complementary-record study—remains open by design.

The empirical contribution is the boundary itself: correction, revalidation, authority, timing, and recovery may belong to one practical governance problem while remaining separately observable in the public record. A human-centered consequential-state regime must ultimately make those relationships legible enough that correction is not merely a change in data, but a governable change in what institutions are still entitled to do with it.

Appendix A. Empirical Control Record

OD-HAS EMPIRICAL CONTROL STATE

Pilot 1 — CFPB
CLOSED — HOLD
25 included records

Ground B: 0 YES / 0 NO / 25 UNKNOWN REBASE-
PENDING: 0 YES / 0 NO / 25 UNKNOWN

Pilot 2 — MSPB
CLOSED — HOLD
20 included chains

Ground B: 1 YES / 19 NO / 0 UNKNOWN REBASE-
PENDING: 0 YES / 1 NO / 19 N/A / 0 UNKNOWN Interim
authority observable: 0 / 20 Resolution clock: 0
COMPLETE / 1 PARTIAL / 19 NONE OBSERVED

Across both pilots:
New Rule A categories 0
New Rule B categories 0
O4 misfits 0
Mid-pilot patches 0
SUT changes 0

- OD-HAS-03 Discovery v0.2
FROZEN — SEARCH SATURATION

Pilot 3:
NO QUALIFYING SINGLE PUBLIC CORPUS IDENTIFIED

Architecture: CLOSED

Appendix B. Claim Boundaries

This paper does not claim: that OD-HAS is universally validated; that the pilot samples estimate institutional prevalence; that a complaint allegation is an adjudicated fact; that Ground B = NO in 19 MSPB chains estimates a general rate; that all public institutions separate basis revision from authority or timing internally; that no qualifying public corpus exists anywhere; that public-record nonobservability establishes institutional absence; that existing law universally requires every OD-HAS disposition; that prospective correction guarantees recovery of residual effects; that the SUT has been independently reproduced in every implementation environment. The paper does claim, more narrowly, that: a frozen cross-domain specification was applied to two heterogeneous real public-record corpora without requiring new Rule A or Rule B categories; Ground B observability changed materially with record structure; CFPB complaint narratives were propagation-rich and procedure-poor; MSPB Board records were grounds-and-reassessment-rich but authority-and-clock-poor;

A directed, frozen discovery program did not identify a single ordinary public corpus that exposed the full Pilot 3 conjunction at the required density; public record architecture is therefore a material condition of construct observability in this research program.

References

- 11 U.S.C. § 362(c)(3).
<https://www.law.cornell.edu/uscode/text/11/362>
- 15 C.F.R. § 766.24.
<https://www.bis.gov/regulations/ear/766>
- 17 C.F.R. § 201.102(e)(3).
- 49 U.S.C. § 44709(e)(3)-(4).
<https://www.law.cornell.edu/uscode/text/49/44709>
- Consumer Financial Protection Bureau. (n.d.). How we share complaint data.
<https://www.consumerfinance.gov/complaint/data-use/>
- Consumer Financial Protection Bureau. (2026a). CFPB Consumer Complaint Database Narratives Archive.
<https://www.consumerfinance.gov/foia-requests/foia-electronic-reading-room/cfpb-consumer-complaint-database-narratives-archive/>
- Consumer Financial Protection Bureau. (2026b, August 14). The CFPB to cease discretionary publication of complaint narratives and visualizations.
<https://www.consumerfinance.gov/about-us/newsroom/the-cfpb-to-cease-discretionary-publication-of-complaint-narratives-and-visualizations/>
- Education Workforce Council. (2026). Fitness to Practise Annual Report for the year ending 31 March 2026.
<https://www.ewc.wales/site/index.php/en/about-us/corporate-information/corporate-plans-and-reports/fitness-to-practise-annual-report-for-the-year-ending-31-march-2026>
- Fair Credit Reporting Act, 15 U.S.C. § 1681i.
<https://www.law.cornell.edu/uscode/text/15/1681i>
- Health and Care Professions Tribunal Service. (n.d.). Search hearings.
<https://www.hcpts-uk.org/hearings/search/>
- Health Professions Order 2001, S.I. 2002/254, art. 31 (UK).
<https://www.hcpc-uk.org/globalassets/resources/legislation/hcpc---consolidated-health-professions-order-2001.pdf>
- Meng, K., Bau, D., Andonian, A., & Belinkov, Y. (2022). Locating and editing factual associations in GPT. *Advances in Neural Information Processing Systems*, 35, 17359-17372. <https://doi.org/10.52202/068431-1262>
- Nursing and Midwifery Council. (2025, September 1). Interim order reviews.
<https://www.nmc.org.uk/ftp-library/interim-orders/interim-order-reviews/>
- Pharmacy Order 2010, S.I. 2010/231, art. 56 (UK).
<https://www.legislation.gov.uk/ukSI/2010/231/article/56>
- Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), 2016 O.J. (L 119) 1. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. *Proceedings of the 2019 Conference on Fairness, Accountability, and Transparency*, 59-68.
<https://doi.org/10.1145/3287560.3287598>
- Shaik, T., Tao, X., Xie, H., Li, L., Zhu, X., & Li, Q. (2025). Exploring the landscape of machine unlearning: A comprehensive survey and taxonomy. *IEEE Transactions on Neural Networks and Learning Systems*, 36(7), 11676-11696. <https://doi.org/10.1109/TNNLS.2024.3486109>
- Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction*, 36(6), 495-504.
<https://doi.org/10.1080/10447318.2020.1741118>
- Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology.
<https://doi.org/10.6028/NIST.AI.100-1>
- Texas State Board of Dental Examiners. (2026). Temporary Suspensions Code.
<https://tsbde.texas.gov/laws-rules/recent-temporary-suspensions/>

U.S. Merit Systems Protection Board. (n.d.-a).
Nonprecedential decisions.
<https://www.mspb.gov/decisions/nonprecdec.htm>

U.S. Merit Systems Protection Board. (n.d.-b). Precedential
decisions. <https://www.mspb.gov/decisions/precdec.htm>

Ustun, B., Spangher, A., & Liu, Y. (2019). Actionable
recourse in linear classification. Proceedings of the
Conference on Fairness, Accountability, and
Transparency, 10-19.
<https://doi.org/10.1145/3287560.3287566>