

Mediated Judgment Under Constraint

Operative Representation, Authority, Burden, and Correction under Finite Action

David T. Swanson

March 18, 2026

Abstract

Judgment in real systems does not proceed over cases in unconstrained fullness. It proceeds through operative representations: files, categories, scores, records, profiles, thresholds, summaries, dashboards, model outputs, and other structured renderings through which cases become actionable for finite agents and institutions. This paper develops *Mediated Judgment Under Constraint*, a framework for explaining how such renderings acquire standing and decisional force, and why the resulting judgment is structurally fallible. Its central claim is that finite action is mediated through operative representations produced under selective constraints and acted upon within judgment regimes. Because such mediation is selective, scope-bounded, and residue-bearing, mediated judgment is vulnerable to patterned misfit not by accident but by structure. Once operative representations govern consequences, their omissions, burden distributions, and reflexive effects become consequential in their own right. The paper argues that traceability, scope discipline, and correction are therefore not optional administrative virtues but constitutive conditions of responsible mediated judgment. Its aim is not to provide a complete ethics or a full political theory, but to specify the structural layer through which finite action becomes authoritative over cases by means of representations that never fully exhaust what they govern.

1 Introduction

1.1. Motivating Problem

Real judgment does not typically meet a case whole.

A hospital does not act on a patient in the full density of lived, clinical, and situational reality. It acts through charts, coded statuses, handoff summaries, orders, and documentation fields. A welfare agency does not act on a person in unconstrained fullness. It acts through forms, eligibility categories, verification records, deadlines, and administratively visible events. A school acts through attendance flags, disciplinary entries, performance indicators, and procedural classifications. A lender, insurer, employer, court, or platform likewise acts through files, profiles, thresholds, scores,

and action-relevant summaries. Such systems do not operate without contact with reality. But neither do they operate over reality in its full depth. They act through renderings of it.

These renderings are not necessarily fictitious. They preserve something real and often indispensable. But they are not identical to the cases they govern. They are selective representations through which finite agents and institutions make cases tractable enough for action. That condition is easy to overlook because the rendering often disappears into the routine of decision. The file comes to stand for the person. The score comes to stand for the trait. The category comes to stand for the condition. The profile becomes the administratively real form of the case for everyone downstream of it.

This matters because once action attaches to a rendering, its omissions cease to be merely descriptive limitations. They become pathways along which decisions travel. A stale record may structure present treatment more authoritatively than present reality. A threshold may convert a gradient into a hard consequence. A score may compress complex circumstance into a ranking that governs allocation. A procedural category may become the form in which a person is recognized, delayed, denied, or redirected. The resulting failures are therefore not well described as errors of information alone. They arise from the structure through which a case becomes actionable in the first place.

This paper begins from the claim that such mediation is not a peripheral defect of institutional or technical systems. It is one of the basic conditions under which finite judgment becomes possible. If that is right, then a theory of judgment cannot stop at rules, reasons, intentions, or outcomes alone. It must also account for the operative representations through which cases become actionable, the judgment regimes through which those representations acquire force, the burdens generated when they misfit the cases they govern, and the correction paths through which such misfit can be surfaced and revised.

1.2. Central Question

The central question of this paper is:

How does finite judgment become authoritative over cases through operative representations, and what follows once such judgment is understood as selective, scope-bounded, and structurally fallible?

1.3. Main Proposal

The paper's main proposal is:

Finite agents and institutions do not act on cases directly, but through operative representations produced under selective constraints and interpreted within judgment regimes. Because such mediation is selective, scope-bounded, and residue-bearing, mediated judgment is structurally fallible. Once operative representations acquire standing and decisional force over cases,

burden visibility, traceability, scope discipline, and correction become constitutive conditions of responsible mediated judgment.

This proposal is not a rejection of procedure, evidence, or practical simplification. It is a claim about their condition. Action under finite circumstances requires selective renderings. The question is therefore not whether judgment can occur without mediation, but how mediated judgment works, where it predictably misfits the cases it governs, and what conditions keep its partiality from hardening into unanswerable authority.

1.4. Roadmap

Section 2 situates the framework against nearby but insufficient ways of framing judgment, representation, and institutional action. Section 3 defines the paper's core objects and distinctions. Section 4 states the main theory of mediated judgment under constraint. Section 5 explains why this theory is needed and what confusions it resolves. Section 6 shows explanatory payoff across institutional judgment, quantification, AI-mediated action, and humanly lived cases. Section 7 addresses the strongest objections. Section 8 clarifies scope, limits, and visible residue. Section 9 sketches implications and future work. Section 10 concludes. Appendix A provides compact provisional definitions, and Appendix B offers a concise structural schema.

2 Background and Rival Framings

2.1. Direct-Judgment Pictures

A common tendency in both ordinary and theoretical discussion is to treat judgment as though it operated on cases directly enough for mediation to remain secondary. Rules are said to be applied to cases. Institutions are said to evaluate persons. Systems are said to classify risk, determine eligibility, or optimize outcomes. This language is not wholly false, but it is structurally compressed. It hides the fact that judgment almost never acts on a case in unconstrained fullness. It acts on some rendering of the case.

This matters because the structure of the rendering shapes the structure of the action. If the rendering is stale, thin, overcompressed, badly scoped, or poorly matched to the stakes of the case, then the resulting judgment may remain orderly while being substantively misaligned. A theory that speaks as if judgment directly meets its object therefore underdescribes the condition under which real action occurs.

2.2. Purely Epistemic Treatments of Representation

A second family of approaches notices mediation but keeps it primarily at the level of epistemology. On such views, representation matters because it affects what can be known, inferred, measured, or justified. This is a genuine insight. But it is incomplete once representations become action-guiding.

A score, file, classification, profile, or model output is not merely an imperfect description when it governs treatment, access, sanction, ranking, or denial. Once it is linked to a decision process, it becomes part of the mechanism through which the case is acted upon. The problem is no longer only what is known. It is what kind of rendering has acquired force over a case and under what conditions that force is exercised.

2.3. Reduction to Ethics Too Early

A third tendency is to move too quickly from representational limitation to ethical evaluation. This move is understandable, because once mediated judgment governs consequences, moral and political stakes appear immediately. But an intermediate layer is still needed. Before one can specify the ethics of mediated judgment, one must explain how judgment becomes mediated, how a rendering becomes operative, how authority attaches to it, how mismatch produces burden, and why correction is structurally necessary rather than merely benevolent.

This paper addresses that middle layer. It does not aim to replace ethics or politics. It aims to clarify the structure that those later theories must reckon with once judgment proceeds through operative representations rather than direct grasp of cases in full.

2.4. The Paper's Position

This paper is written to stand on its own. It does not require the reader to accept a larger theoretical canon in order to understand its argument. Its claims are meant to be intelligible as a self-standing account of mediated judgment under finite action. At the same time, the framework is compatible with broader concerns about finite description, selective rendering, and the limits of action through partial representations. That compatibility matters architecturally, but it is not a condition of local intelligibility.

3 Core Definitions and Distinctions

This section introduces the paper's central objects and the distinctions that organize them. The aim is not terminological multiplication for its own sake. The aim is to separate layers that are often collapsed in both ordinary and scholarly discussion. When the same language is used for the case, the traces through which it becomes available, the rendering through which it becomes actionable, the regime that acts over that rendering, and the costs produced when the rendering misfits the case, failures become harder to locate and harder to correct. The framework developed here therefore distinguishes case from operative representation, trace from evidence, standing from decisional force, constraint from burden, and correction from discretionary rescue.

3.1. Case and World-Fragment

A *case* is the practical target of judgment: the person, situation, claim, event, process, or problem toward which action is directed. The ordinary term is useful because this paper is concerned not only with abstract representation but with real judgment over concrete matters.

A *world-fragment* is the more formal counterpart to *case*. The term is retained because not every target of judgment is a simple, already-bounded object. A world-fragment may be a patient in treatment, a welfare claim, a student's institutional status, a credit applicant, a risk-bearing population, or a compound practical circumstance requiring coordinated response. In the main prose, *case* will usually be preferred for readability, but the underlying point is the same: what judgment acts toward is not identical to the rendering through which it becomes tractable.

3.2. Operative Representation

An *operative representation* is the rendering through which judgment becomes action-guiding. This includes files, profiles, categories, scores, metrics, thresholds, case summaries, coded records, dashboards, model outputs, and other structured renderings stable enough for procedure, communication, or intervention. [1, 8]

The term is meant to be broad without becoming indiscriminate. Not every representation is operative in the relevant sense. A representation is *operative* when consequences for a case are materially routed through it within a judgment process. A rendering may exist, matter, or influence perception without yet functioning as the object over which consequences are organized. An operative representation is one that has entered the path of judgment in a way that helps determine what follows.

3.3. Access Trace

An *access trace* is a measurement, report, testimony, event record, perceptual uptake, document, or other trace through which some aspect of a case becomes available. The concept is meant to remain closer to access than to settled evidentiary force. A trace is something through which the case becomes available at all, not yet something that automatically counts in support of a judgment.

This concept matters because systems often move too quickly from what is available to what is taken to matter. The paper therefore needs a term for the earlier layer at which some feature of a case becomes accessible without yet being action-relevant.

3.4. Evidentiary Interpretation

An *evidentiary interpretation* is an access trace taken under a standard of relevance, admissibility, credibility, or sufficiency. The point of this term is to mark that evidence is not simply found. A trace becomes evidence only within some regime that specifies what it is taken to show, how strongly it counts, and whether it counts at all.

The same trace may function as strong evidence in one setting, weak evidence in another, and irrelevant material in a third. This distinction matters because mediated judgment depends not only on what became available, but on what a system was prepared to treat as support for decision.

3.5. Conditioned Cut

A *conditioned cut* is the structured selection-regime through which a case is rendered into a usable representation under finite conditions. It determines what distinctions are preserved, what is backgrounded, what format the case must take, and what becomes tractable enough for judgment or action.

A cut is *conditioned* because it is shaped by constraints, tools, procedures, institutional formats, purposes, and practical stakes. The concept matters because it blocks the idea that simplification is merely an accidental defect. Selective rendering is part of what makes mediated judgment possible at all.

3.6. Judgment Regime

A *judgment regime* is the rule, procedure, institutional logic, workflow, or decisional process that acts over an operative representation. The regime specifies how the rendering is interpreted, what thresholds matter, how outputs are translated into consequences, and what kinds of override, review, or exception are available.

This concept matters because an operative representation does not govern consequences on its own. Its practical force depends on the regime acting over it.

3.7. Standing Condition and Decisional Force

A *standing condition* is the condition under which a representation is allowed to enter judgment as action-relevant. It marks what gives a rendering enough recognized status to be treated as admissible, usable, or institutionally legible within a judgment regime.

Decisional force is the degree to which an operative representation actually governs what follows in a case. A representation may have standing without being outcome-determining. It may count in the process without materially routing the result. Decisional force marks the stronger condition in which consequences are in fact organized through the rendering.

This split matters because it separates two questions that are often run together: what is allowed to count, and what actually governs outcome.

3.8. Scope Condition

A *scope condition* marks the range within which an operative representation remains adequately action-guiding. No rendering is equally fit for all purposes, stakes, scales, or use-contexts. Scope

therefore identifies where a representation's authority remains defensible and where it begins to degrade.

This concept matters because many failures are not failures of total falsehood. They are failures of travel beyond legitimate use. A representation may be useful, orderly, and even well-supported for one task while becoming misleading when exported to another task, scale, or consequence context.

3.9. Residue

Residue is the structured remainder relative to a rendering: what the operative representation omits, flattens, merges, backgrounds, or carries badly. Residue is not a mystical outside. It is the ordinary remainder produced whenever finite judgment depends on selective stabilization.

At a minimum, residue may take recurrent forms such as omitted contextual structure, merged or flattened distinctions, and omitted lived or burden-relevant features of the case. The concept matters because a useful rendering is easily mistaken for an exhaustive one. Residue marks what the rendering leaves behind even when it works.

3.10. Constraint and Burden

Constraint names unavoidable limits on time, memory, staffing, computation, money, coordination, training, and interpretive bandwidth. Constraints help explain why mediation and simplification are necessary conditions of finite judgment.

Burden names the costs imposed on those who must live under, navigate, satisfy, or correct mediated judgment. These burdens may be documentary, procedural, temporal, interpretive, emotional, or logistical. Constraints are unavoidable; burdens are distributed. [6, 2]

This distinction matters because institutions often treat their own limits as though they automatically justified the burdens their systems impose. The theory needs to separate the fact of finitude from the question of who absorbs the costs of simplification.

3.11. Traceability

Traceability is the degree to which a mediated judgment leaves inspectable traces of how a case was rendered, what counted, how the judgment regime operated, and where the result came from. Traceability does not require perfect transparency. It requires enough visibility into the path from case to consequence that challenge, review, and diagnosis remain genuinely possible.

This concept matters because correction depends on more than the abstract possibility of appeal. If the route from case to representation to consequence is opaque, correction becomes theatrical rather than usable.

3.12. Correction Path

A *correction path* is a pathway for challenge, amendment, override, appeal, rollback, recalibration, audit, exception-handling, second review, or redesign. It is useful to distinguish two levels:

- **case-level correction**, such as appeal, override, amendment, or exception in a particular case;
- **system-level correction**, such as audit, recalibration, redesign, retraining, or rule revision aimed at patterned failure.

Correction matters because once mediated judgment is structurally fallible, revision cannot be treated as an optional courtesy. It becomes part of what responsible mediated judgment requires.

3.13. Structural Fallibility

Structural fallibility is the standing possibility of patterned misfit arising from finite action through selective operative representations. It does not mean that judgment always fails, nor that all representations are equally unreliable. It means that mismatch is built into the condition of mediated action itself and therefore cannot be treated as a rare anomaly.

Structural fallibility may take the form of stale records, proxy drift, threshold distortion, omitted context, misclassification, forced commensuration, or cross-scale misfit. What unifies these phenomena is not one mechanism of error, but the fact that finite action proceeds through selective renderings rather than unconstrained contact with the case in full.

3.14. Reflexive Consequence

Reflexive consequence refers to the way mediated judgment can reshape the very case it governs. Once a rendering is tied to action, it may alter treatment, opportunities, institutional status, self-understanding, or future traces. The representation does not merely describe the case; it enters the case's ongoing organization. [4]

This concept is extension-level rather than strictly required in every application. It becomes especially important where judgments feed back into the conditions they purport only to register.

3.15. Load-Bearing Distinctions

Several distinctions organize the rest of the paper.

The first is the distinction between *case* and *operative representation*. This blocks the recurrent confusion in which files, scores, profiles, or records are treated as if they exhausted the cases they govern.

The second is the distinction between *access trace* and *evidentiary interpretation*. This blocks the confusion between what became available and what a system was prepared to count.

The third is the distinction between *operative representation* and *judgment regime*. This blocks

the compression in which all failure is attributed either to bad data or bad judgment without any account of the structure that links them.

The fourth is the distinction between *standing condition* and *decisional force*. This blocks the confusion between what is allowed to count and what actually governs outcome.

The fifth is the distinction between *constraint* and *burden*. This blocks the inference that institutional finitude automatically justifies the costs shifted onto subjects.

The sixth is the distinction between *scope failure* and *total falsehood*. This blocks the crude alternative between perfect adequacy and complete uselessness.

The seventh is the distinction between *correction path* and *discretionary rescue*. This blocks the idea that revision is merely exceptional kindness rather than a structural response to fallibility.

The eighth is the distinction between *traceability* and mere *formal reviewability*. This blocks the appearance of accountability where the actual route from case to consequence remains too opaque for meaningful challenge.

4 Main Theory of Mediated Judgment Under Constraint

4.1. Field-Level Characterization

Mediated Judgment Under Constraint is a theory of how finite action becomes authoritative over cases through representations that never fully exhaust what they govern. It is not a general metaphysics, not merely an epistemology of representation, and not yet a full ethics. Its object is narrower and more structural: the middle layer through which cases are rendered operable, admitted into judgment, acted upon under a regime, and exposed to the consequences of selective mediation.

The theory addresses a recurring practical condition. Real agents and institutions do not decide over cases in unconstrained fullness. They decide through renderings of those cases. Those renderings are necessary for action, but they are also selective, scope-bounded, and vulnerable to mismatch. The present theory explains that condition, the forms of failure to which it gives rise, and the structural reasons burden and correction become central once such renderings acquire force over outcomes.

4.2. The Mediated-Judgment Premise

The central premise of the theory is that finite agents and institutions do not act on cases directly, but through operative representations. These representations may be formal or informal, quantified or narrative, technical or administrative. What unifies them is not surface form but practical role: they are the renderings through which consequences are materially routed.

This premise should be read carefully. It does *not* imply that action is severed from reality. The case remains what the representation is about, and it remains the source of possible correction precisely because it exceeds the rendering through which it has been made actionable. The point is narrower and more disciplined: finite judgment can proceed only by stabilizing some representation

of the case strongly enough for a judgment regime to work over it.

4.3. Why Mediation Is Unavoidable

Mediation is unavoidable because finite action requires selective compression. No agent or institution can preserve every distinction equally live at once. Action requires some combination of selection, grouping, formatting, standardization, and reduction. A rendering that preserved every feature without loss would cease to function as a practical object for judgment.

This becomes especially clear at institutional scale. Large systems must coordinate across time, offices, personnel, procedures, and populations. They therefore require representations that travel: files, forms, categories, records, profiles, scores, dashboards, flags, and standardized summaries. That portability is necessary. But it always comes at a cost. What travels well is thinner than the lived, local, or situationally rich reality it represents.

Mediation should therefore not be treated as a regrettable supplement added to otherwise transparent access. It is an enabling condition of finite judgment itself.

4.4. From Representation to Consequence

Not every representation has the same practical significance. A rendering becomes consequential only when it becomes operative within a judgment regime. This requires more than mere presence. It requires, at minimum, that the representation satisfy a *standing condition*—that is, that it count as admissible or action-relevant within the process—and that it acquire enough *decisional force* to help determine what follows.

This distinction matters. A representation may have standing without materially routing outcome. It may count in the file, appear in the workflow, or be visible to the system while remaining weak in effect. By contrast, a representation has decisional force when outcomes are in fact organized through it, absent override or correction.

Mediation becomes practically consequential when a rendering crosses that threshold from relevance to force. At that point, the representation no longer serves only as a rough guide. It becomes part of the mechanism through which the case is handled. A threshold may divide eligibility from denial. A score may determine ranking. A profile may trigger intervention. A coded record may stand in for a person's present condition more authoritatively than the person's own report. A dashboard flag may command attention while uncoded reality remains inert. In each case, the representation does not replace the case, but it becomes the form through which the case is practically governed.

4.5. Structural Fallibility

Because operative representations are selective, scope-bounded, and residue-bearing, mediated judgment is structurally fallible. Fallibility here does not mean that judgment always fails, nor that all representations are equally unreliable. It means that mismatch is a standing possibility internal to the structure of mediated action itself.

This mismatch may take many forms: stale records, omitted variables, proxy drift, threshold distortion, misclassification, forced commensuration, cross-scale misfit, or failure to preserve what matters for the relevant stakes. What unifies these failures is not one mechanism of error, but the fact that finite action proceeds through selective renderings rather than unconstrained access to cases in full.

The key point is therefore not merely that systems sometimes make mistakes. It is that error cannot be treated as an occasional anomaly layered onto an otherwise exhaustive relation between representation and case. Once judgment depends on selective stabilization, patterned misfit becomes structurally possible.

4.6. Scope and Overextension

Structural fallibility becomes especially important because operative representations do not fail only by being false. They also fail by traveling beyond the conditions under which they remain adequately action-guiding. This is the problem of scope.

A representation may be adequate for one task, one scale, one level of consequence, or one decision environment, while becoming misleading in another. A score may be useful for rough triage while being unfit for punitive ranking. A file may be adequate for coordination while being inadequate as a morally or practically exhaustive rendering of the case. A threshold may be administratively useful while becoming distortive when treated as if it captured a natural boundary in the world.

For this reason, many failures of mediated judgment are better understood as *overextensions* than as total falsehoods. The representation retains decisional force after the conditions of its adequate use have broken down.

4.7. Burden Distribution

When authoritative renderings misfit the cases they govern, the resulting reconciliation costs do not disappear. They are distributed. Often they are shifted downward onto those least able to absorb them: the claimant who must gather more documentation, the patient who must restate what the chart did not carry, the applicant who must infer hidden criteria, the citizen who must navigate delay, repetition, ambiguity, or procedural opacity.

This is why the distinction between constraint and burden matters so much. Institutions may indeed face real constraints. But those constraints do not determine who should absorb the costs of simplification. Burden distribution is one of the main ways mediated judgment becomes socially and politically consequential. A system may present itself as merely operating under finite limits while, in practice, exporting the labor of repair, translation, and reconciliation onto those subject to its authority.

4.8. Correction as Structural Condition

If fallibility is structural, then correction must also be structural. A system that acts through selective renderings while refusing meaningful paths for challenge, amendment, override, recalibration, or redesign behaves as though its operative representations were self-exhausting. That is not a neutral stance. It is a denial of the very conditions under which mediated judgment operates.

Correction should therefore not be treated as a discretionary kindness added after judgment has already done its real work. It is part of what keeps mediated judgment answerable to the cases it governs.

It is useful to distinguish two levels here. *Case-level correction* includes appeal, amendment, override, exception, or second review in a particular case. *System-level correction* includes audit, recalibration, redesign, retraining, or rule revision aimed at patterned failure. Both matter. A system that allows only case-level rescue may preserve local mercy while leaving structural distortion intact. A system that allows only system-level review may acknowledge abstract error while leaving concrete cases trapped in the meantime.

4.9. Traceability

Correction depends on traceability. A judgment process must leave enough inspectable traces of how the case was rendered, what counted, how the judgment regime operated, and where the result came from for challenge to be more than ceremonial.

Traceability does not require perfect transparency. It requires enough visibility into the path from case to operative representation to consequence that review, diagnosis, and revision remain genuinely possible. Where that path is opaque, correction becomes theatrical. The formal existence of appeal or review then fails to amount to practical answerability.

4.10. Reflexive Consequence

In many domains mediated judgment does not merely act on a pre-existing case. It helps reshape the case. A diagnosis may alter treatment and self-understanding. A school classification may redirect a student's path. A risk score may affect surveillance intensity. A welfare file may reorganize a person's practical options. Once tied to action, the mediated rendering does not simply describe the case; it enters the case's ongoing organization.

This reflexive consequence is not equally central in all domains. But where it arises, it sharpens the stakes of mediated judgment considerably. The costs of mismatch do not remain external to the case. They may become part of what the case now is.

4.11. Compact Structural Schema

The theory can be represented schematically as:

$$\begin{aligned}
 W &\xrightarrow{R_f} R + \Delta \\
 \Sigma(R, J) = 1 \wedge \Phi(R, J) &> 0 \\
 J(R \mid S, \Sigma, \Phi) &\rightarrow A \\
 A &\rightarrow B \\
 A, R &\rightarrow W' \\
 K(A, R, \Delta, B, J) &\rightarrow \text{revision}
 \end{aligned}$$

where:

- W = case or world-fragment,
- R_f = rendering formation under finite constraint,
- R = operative representation,
- Δ = residue,
- J = judgment regime,
- S = scope conditions,
- Σ = standing condition,
- Φ = decisional force,
- A = action,
- B = burden distribution,
- W' = the case as potentially reshaped by mediated action,
- K = correction path.

The point of the schema is not theorematic closure. It is structural clarity. A case is rendered under finite conditions; that rendering leaves residue; it acquires standing and decisional force within a regime; action is then taken through it; burdens follow when mismatch is exported rather than absorbed; the case may be reshaped by that action; and correction is required if the process is to remain answerable to what it governs.

5 Why This Theory Is Needed

5.1. The Missing Structural Layer

This theory is needed because many existing discussions move too quickly across levels that should be kept distinct. They pass from case to trace, from trace to evidence, from evidence to representation, from representation to decision, and from decision to outcome without making the intervening structure fully explicit. When that happens, failure becomes harder to locate and easier to misdescribe. A bad outcome may be blamed on a bad rule when the operative representation was already inadequate. A bad representation may be blamed when the sharper source of distortion lay in the judgment regime acting over it. Institutional burden may be treated as an unfortunate side effect of necessity rather than as a distributive result of how mediated authority is organized.

What is missing in such discussions is not simply more detail. It is a distinct structural layer. A theory of mediated judgment makes explicit the sequence through which cases become actionable: what becomes available, what is interpreted as evidence, how a rendering acquires standing, how it gains decisional force, and how the resulting action distributes costs when the rendering misfits what it governs. Without that layer, critique tends to remain coarse. One can say that something went wrong, but not with sufficient precision where it went wrong or why the error took the form it did.

5.2. What Confusions It Resolves

The framework resolves several recurring confusions.

First, it resolves the confusion between the case and its rendering. A file is not the person. A score is not the trait. A profile is not the life. A record is not the situation in full. This matters because institutions routinely act as though the administratively portable rendering had become the case itself.

Second, it resolves the confusion between what was available and what counted. Not every access trace becomes evidence in the same way. A report may exist without being admissible. A document may be present without carrying weight. A testimony may be available while remaining structurally discounted. The theory therefore distinguishes what became accessible from what a regime was prepared to treat as action-relevant support.

Third, it resolves the confusion between representation and decision. Operative representations and judgment regimes are not the same object. Each can fail differently. A rendering may be thin, stale, or misaligned even if the decision rule is internally consistent. A decision rule may be crude, severe, or badly scoped even if the rendering it acts over is relatively rich. Keeping these layers distinct makes diagnosis more precise.

Fourth, it resolves the confusion between institutional constraint and justified burden. Real systems face limits. They cannot process cases in unconstrained fullness. But that fact does not determine who should absorb the costs of simplification. The theory shows that burden is not simply the shadow cast by finitude. It is also a distributive output of how mediated authority is designed and exercised.

Fifth, it resolves the confusion between procedural order and adequacy. A system may be stable,

repeatable, auditable, and formally orderly while still remaining misaligned with the cases it governs. Orderliness is not nothing, but it is not by itself a guarantee that the representation is fit for the consequence structure attached to it.

Sixth, it resolves the confusion between correction and discretionary rescue. Revision is not merely kindness after the fact. It follows from the structure of mediated judgment itself. If action proceeds through selective, residue-bearing renderings, and if those renderings govern outcomes, then challenge, amendment, override, and redesign are not optional supplements. They are part of what keeps judgment answerable to the cases it governs.

5.3. Why the Structural Gain Is Real

The gain of the theory lies in making explicit a layer that is usually implicit but constantly operative. The framework shows:

1. how finite judgment requires selective renderings,
2. how some renderings become operative rather than merely descriptive,
3. how operativity becomes standing and decisional force within a judgment regime,
4. why authority exercised through such renderings is structurally fallible,
5. how mismatch becomes distributed burden rather than disappearing,
6. and why traceability and correction must be built into responsible mediated judgment rather than appended after it.

That is a real gain because it gives later ethical, political, and design arguments a clearer object. Instead of speaking vaguely about institutions, rules, or systems in general, one can ask more disciplined questions: What representation was operative? Under what standing condition? With what decisional force? Under what scope condition? What residue became relevant? Who bore the reconciliation costs? What correction path remained available? Once those questions are available, the structure of failure becomes more legible, and the structure of possible repair becomes more concrete.

6 Applications and Explanatory Payoff

The framework developed here is meant to clarify real domains of judgment rather than remain a purely internal architecture. Its explanatory payoff lies in showing how the same structural pattern recurs across otherwise different settings: a case is selectively rendered into an operative representation, that representation acquires standing and decisional force within a judgment regime, and the resulting action distributes burdens when the rendering misfits what it governs. The following applications are not exhaustive. They are meant to show the theory doing concrete work.

6.1. Institutional Judgment

Institutional systems routinely act through operative representations. Hospitals act through charts, coded statuses, handoff summaries, and order sets. Welfare agencies act through forms, deadlines, verification states, and administratively visible events. Schools act through attendance records, classifications, behavioral entries, and performance indicators. Courts act through legally admissible renderings of cases rather than through the case in unconstrained fullness. In each setting, judgment becomes portable only by becoming selective.

The framework clarifies why such systems so often generate friction without requiring overt malice or obvious bad faith. Their operative representations are necessary, but they are also partial. Once consequences attach, what they fail to carry becomes institutionally consequential. A hospital chart may preserve medication history while failing to preserve what a condition is like for the patient living through it. A welfare file may preserve documentary status while failing to preserve the practical labor required to remain legible to the system. A school classification may preserve administratively visible conduct while failing to preserve the circumstances through which the conduct became intelligible. The problem, in each case, is not that representation exists. It is that a selective rendering acquires force over a case while the costs of its selectivity are displaced elsewhere. [1, 3, 6]

A simple welfare-style case helps show the structure more explicitly. A claimant's situation first appears through access traces: income records, residency documents, employment history, household status, deadlines met or missed, and prior administrative entries. These traces are then interpreted under the evidentiary rules of the system. The operative representation becomes the administratively legible case: verified, pending, noncompliant, incomplete, or eligible. Once that rendering acquires standing and decisional force, the burden of mismatch falls somewhere. If the representation fails to carry irregular work history, unstable housing, or the practical reasons for missing a deadline, the system's simplification costs do not disappear. They are pushed onto the claimant as documentary, temporal, interpretive, and emotional burden. Correction then depends on whether a usable path of appeal, amendment, or review still exists.

6.2. Quantification

Quantified judgment is especially revealing because its outputs often appear impersonal, precise, and objective. Scores, indicators, rankings, thresholds, and performance measures can be extremely useful. They compress complexity into forms that travel well across institutions, offices, and timescales. But they remain operative renderings rather than exhaustive captures. Their authority often exceeds their scope, especially when numerical neatness is mistaken for adequacy.

The framework helps explain why the problem is not numbers as such. The problem is the authority granted to compressed renderings without sufficient scope discipline, burden visibility, or correction. A threshold can make action possible while still imposing a hard division on what is in reality a gradient. A score can enable comparison while still carrying forward proxy distortions or hidden assumptions about what matters. A ranking can appear fair because it is formally consistent while

remaining misaligned with the underlying case structure. Quantification therefore does not eliminate mediated judgment. It intensifies it by stabilizing a rendering in especially portable form. [1, 8, 7]

The point is especially clear when quantification travels beyond the task for which it was built. A risk score designed for rough triage may later be used for punitive ranking. A performance metric designed for aggregate oversight may become a tool for individual discipline. A cost proxy may be treated as if it tracked need itself. In such cases the representation is not simply wrong everywhere. It is overextended. Its decisional force survives after the conditions of its adequate use have broken down.

6.3. AI and Automated Judgment

AI systems make the structure of mediated judgment unusually visible. Inputs are already selective renderings of cases. Models transform those renderings into outputs such as classifications, predictions, recommendations, summaries, or rankings. Those outputs then enter larger judgment regimes in which they may acquire standing and decisional force. The resulting systems are therefore not wholly novel in structure. They are intensified instances of a more general condition: finite action through selective operative representations.

The framework helps explain why AI raises urgent questions of traceability, burden, proxy drift, scope, and revision without requiring the claim that AI alone created these problems. A model may be formally impressive while still depending on input features that carry poor proxies for what actually matters. A system may appear consistent while still exporting the labor of correction onto those who must contest its outputs. A classification may be technically legible while remaining practically unreadable to the person governed by it. What AI changes, in many cases, is not the existence of mediated judgment but its scale, opacity, speed, and rigidity. [8, 2, 7]

This also clarifies why traceability matters so much in automated settings. When a model output acquires standing and decisional force inside a larger workflow, the practical question is no longer only whether the model is accurate on average. It is also whether the route from case to output to consequence remains inspectable enough for challenge and revision. Where that path is too opaque, correction becomes formal rather than usable.

6.4. Humanly Lived Cases

Where the case being governed is humanly lived, the stakes often intensify. An operative representation may preserve administratively or technically useful structure while failing to carry burden, fear, humiliation, practical unreadability, or lived disorientation. A case may be legible to the system in one register while remaining badly represented in another. This matters because the omissions of mediated judgment may become consequential not only structurally but experientially and practically.

The present paper does not attempt a full adequacy theory of lived cases. Its narrower claim is that where the target is a human being or other lived situation, mediated judgment often encounters a sharper version of its own general problem. The rendering may preserve enough to guide action

while still failing to preserve what it is like to undergo the case from within. Such failures are not always ethically decisive on their own, but they are rarely inert. They shape what kinds of burden, opacity, and correction become possible. [3, 5, 4]

This is one reason humanly lived cases should not be treated as merely more complicated instances of administrable objects. They are cases in which omitted context may also be omitted significance. A system may remain procedurally coherent while still failing to register what its own handling of the case is like for the person being handled.

6.5. Ordinary Human Judgment

The framework also applies below the level of formal institutions. Human beings judge each other through compressed narratives, remembered fragments, selective attention, role categories, and practical shortcuts. Everyday judgment does not therefore become impossible or illegitimate. But it does mean that revisability, scope discipline, and resistance to premature totalization are not only institutional virtues. They are conditions of responsible judgment more generally.

This extension matters because it shows that mediated judgment is not a special defect of bureaucracy or technology alone. Institutional and technical systems intensify the structure, formalize it, and attach larger consequence chains to it, but they do not create it from nothing. The broader lesson is that judgment in general proceeds through renderings rather than through unconstrained access to cases in full. Formal systems merely make this condition harder to ignore. [5, 9]

7 Objections and Replies

7.1. Objection 1: This Is Just Epistemology

One might object that the framework is nothing more than an epistemology of limited access stated in institutional language. On that reading, the paper would amount to the familiar claim that finite agents do not know everything and must therefore act under uncertainty.

That objection misses the paper's central shift. The framework does begin from a representational condition, but it does not remain at the level of knowledge alone. Once a rendering becomes operative within a judgment regime, its limits no longer remain merely epistemic. They become part of the structure through which consequences are imposed. A file, score, category, or model output is not just an imperfect description sitting beside the case. It may become the object through which the case is handled, delayed, denied, ranked, or redirected. The theory therefore concerns not only what can be known, but how selective renderings acquire standing and decisional force over cases. Its object is mediated judgment, not epistemic limitation in abstraction from action.

7.2. Objection 2: This Collapses into Skepticism

A second objection is that if all mediated judgment is partial, selective, and residue-bearing, then the framework collapses into generalized distrust. If every rendering leaves something out, perhaps no judgment can claim serious authority.

That conclusion does not follow. Structural fallibility is not the claim that judgment always fails, nor that all representations are equally unreliable. It is the claim that judgment proceeds through selective renderings and is therefore vulnerable to patterned misfit. The point of the framework is not to abolish action, but to discipline it. A representation may still be highly useful, well-supported, and locally adequate. What the theory rejects is the stronger inference from practical success to exhaustive authority. Partiality does not imply paralysis. It implies that judgment must remain answerable to the conditions under which it operates.

7.3. Objection 3: This Explains Too Much

A third objection is that the framework risks becoming too broad. Once one notices mediated representation everywhere, it becomes tempting to redescribe every failure of judgment as a mediation problem. At that point, the theory would explain everything too quickly and therefore explain very little well.

This objection should be taken seriously. The reply is that the paper does not claim that all failures are reducible to mediation. It claims only that mediated representation is one structurally underdescribed layer of judgment and that many recurring failures cannot be well located without it. Some failures arise from corrupt motives, unjust rules, political incentives, material scarcity, or institutional cruelty. The present framework does not replace those explanations. It identifies an additional layer that often shapes how such failures become actionable. Its burden is therefore architectural rather than totalizing.

7.4. Objection 4: Operative Representation Is Too Broad

A fourth objection is that the concept of operative representation may be too broad to do useful work. If every file, score, summary, category, and model output counts as operative, the term risks becoming a catch-all rather than a disciplined object.

This is a genuine pressure point. The framework answers it by tying operativity not to mere presence, but to action-guiding force within a judgment regime. A representation is operative in the relevant sense only when consequences for a case are materially routed through it. A rendering may exist in the background, influence perception, or appear somewhere in a file without yet being operative. The term is therefore not meant to cover every representation equally. It names those renderings that actually enter the path by which judgment becomes consequential.

7.5. Objection 5: Correction Is Being Smuggled In

A fifth objection is that the paper simply assumes correction as a moral ideal. On this view, the argument would not derive correction from its own structure, but would import it from an independent preference for revisability, flexibility, or procedural fairness.

The reply is that correction is introduced here first as a structural requirement rather than as a free-floating moral aspiration. If judgment is mediated through selective renderings that do not exhaust the cases they govern, and if those renderings nonetheless acquire standing and decisional force, then misfit is not an accidental oddity. It is a standing possibility internal to the architecture of mediated action. Once that is granted, pathways of revision are not decorative additions. They are part of what keeps the system answerable to the cases it governs. The paper does not yet claim that every form of correction is equally just, nor that revision alone guarantees legitimacy. Its narrower claim is that a system with no usable correction path behaves as though its own operative renderings were self-exhausting. That is a structural, not merely ornamental, point.

7.6. Objection 6: Burden Belongs to Ethics or Politics, Not Here

A final objection is that burden distribution belongs properly to moral or political theory rather than to a structural account of judgment. On this view, the present paper would overstep its domain by treating burden as one of its own core outputs.

The reply is that the paper does not yet offer a full moral evaluation of burden, nor a complete politics of its distribution. It identifies burden as one of the ways representational mismatch becomes consequential once mediated authority governs cases. That diagnostic step belongs here even if its fuller evaluation belongs partly elsewhere. If a system acts through selective renderings and exports the labor of repair, translation, delay, or legibility downward onto those subject to its authority, then burden is not an external aftereffect. It is one of the forms taken by mediated misfit. Later work may ask whether such burdens are fair, just, or legitimate. The present paper's more limited task is to show why they arise structurally in the first place.

8 Scope, Limits, and Residue

8.1. Scope Conditions

This theory applies wherever action is guided through selective renderings of cases under finite conditions. Its strongest and most immediate relevance is to institutional, quantified, and technically mediated judgment, where operative representations acquire enough standing and decisional force to organize consequences in stable ways. But the framework is not limited to those domains. It is meant to describe a more general structural condition of mediated judgment, even where the forms of representation, review, and correction are less formalized.

8.2. What the Paper Does Not Claim

The paper does not claim:

- that all judgment reduces to formal modeling or explicit data structures,
- that all mediated judgment is equally defective or equally risky,
- that correction by itself guarantees justice,
- that procedural review is sufficient for legitimacy,
- that burden is exhausted by the diagnostic categories introduced here,
- or that the present framework already amounts to a complete ethics or political theory.

Its ambition is narrower. The paper isolates the structural layer through which selective renderings become action-guiding over cases and shows why that layer generates recurring problems of misfit, burden, and revision.

8.3. Open Problems

Several problems remain open.

First, the exact criterion of operativity still needs sharpening. The paper argues that a representation is operative when consequences are materially routed through it, but a more exact account of that threshold remains to be developed.

Second, the distinction between standing and decisional force is structurally useful but not yet fully formalized. The theory can identify the difference between what is allowed to count and what actually governs outcome, but it does not yet offer a complete account of how these conditions vary across judgment regimes.

Third, burden distribution requires more explicit treatment if it is to support later normative work. The present paper shows why burdens arise structurally, but not yet how they should be comparatively assessed, weighted, or politically evaluated.

Fourth, the theory of correction is stronger in architectural justification than in procedural detail. The paper explains why correction must exist, but it does not yet specify the governance conditions under which particular forms of appeal, audit, override, or redesign become adequate.

Fifth, reflexive consequence remains in need of sharper treatment. The paper shows why mediated judgment may reshape the cases it governs, but it does not yet give a full account of when such feedback is central, when it is secondary, and how its force should be modeled.

8.4. Residues of the Theory Itself

This theory leaves visible residue of its own.

It is stronger on structural diagnosis than on formal specification. It explains why mediated judgment

requires its own layer, but it does not yet fully formalize how judgment regimes should be modeled across domains. It shows why correction is structurally necessary, but it does not yet specify the full governance conditions of legitimate review. It clarifies why burden is one of the main ways mismatch becomes consequential, but it does not yet provide a complete account of how different burdens interact or how they should be evaluated downstream.

These are not defects to be concealed. They are part of the theory's present boundary. A framework concerned with selective rendering, structural fallibility, and answerability should make its own unfinished regions visible rather than pretending to have eliminated them. The theory's current contribution is therefore not final closure, but a clearer structural map of a layer that is often active in practice while remaining underdescribed in theory.

9 Implications and Future Work

9.1. Architectural Implication

The main architectural implication of the paper is that theories of judgment should separate layers that are often collapsed into one another. At minimum, four layers must be kept distinct:

1. the case or world-fragment toward which judgment is directed,
2. the operative representation through which that case becomes actionable,
3. the judgment regime acting over that representation,
4. and the burdens and correction demands that arise once consequences attach.

This layered view matters because each level can fail differently. A case may be badly rendered, a rendering may be acted on under a crude regime, a regime may overextend the authority of a representation, and the resulting mismatch may be exported downward as burden rather than absorbed through revision. Once these layers are separated, both diagnosis and repair become more precise.

9.2. Implications for Real Systems

Any real system acting through operative representations should therefore be evaluated not only by procedural order, formal regularity, or predictive performance. It should also be evaluated by the quality of the rendering through which it acts, the scope conditions under which that rendering remains adequately action-guiding, the burdens generated when it misfits the case, the degree of traceability preserved in the path from case to consequence, and the correction paths left available when the system gets things wrong.

This shifts the evaluative focus. The relevant question is not only whether a system reaches outcomes efficiently or consistently. It is also whether the route by which it reaches them remains answerable to the cases it governs.

9.3. Future Work

Several directions for future work follow from this framework.

One direction is normative. The present paper isolates the structural layer through which selective renderings become action-guiding. A later step would be to develop a fuller ethics of mediated action built on that account.

A second direction is institutional and design-oriented. If mediated judgment is structurally fallible, then institutions need better ways of specifying scope, preserving traceability, distributing burdens more responsibly, and building usable correction into ordinary practice rather than treating it as exceptional rescue.

A third direction concerns AI governance and documentation. The framework suggests that technical systems should be examined not only for accuracy or fairness in the abstract, but also for the standing and decisional force of their outputs, the burden of contesting them, and the forms of correction their deployment actually permits.

A fourth direction concerns administrative burden. The present paper shows why burden is structurally generated once selective renderings govern consequences. Future work can ask how such burdens should be measured, compared, and normatively assessed across domains.

A fifth direction is formal. The theory would benefit from a more explicit treatment of judgment regimes, standing conditions, decisional force, scope breach, and correction dynamics. Such formalization would not replace the framework's conceptual work, but could sharpen its use in diagnosis and system design.

10 Conclusion

10.1. Main Result

The main result of this paper is a self-standing theory of mediated judgment under finite constraint. The paper has argued that finite agents and institutions do not act on cases in unconstrained fullness, but through operative representations that acquire standing and decisional force within judgment regimes. Because such representations are selective, scope-bounded, and residue-bearing, mediated judgment is structurally fallible. Once these renderings govern consequences, burden distribution, traceability, and correction become central rather than peripheral features of responsible judgment.

10.2. Broader Significance

This matters because many failures of judgment do not arise only from bad motives, bad rules, or bad ends. They also arise because selective renderings acquire force without adequate scope discipline, without sufficient visibility into the burdens they export, and without usable paths of revision when they misfit the cases they govern. The paper therefore clarifies an underdescribed layer of real action: the structure through which selective representation becomes operative authority

over cases.

That clarification has broader value because it gives later ethical, political, and design arguments a more precise object. It becomes possible to ask not only whether a decision was fair, efficient, or lawful, but also what representation was operative, how it acquired force, what it left out, who absorbed the costs of that omission, and what correction remained available once its authority was exercised.

10.3. Final Claim

Mediated Judgment Under Constraint is the view that finite agents and institutions act through operative representations rather than over cases in unconstrained fullness, and that once such representations acquire standing and decisional force, their selectivity, burden distributions, and need for correction become structurally central.

If that is right, then the problem is not how to act without mediation. It is how to act responsibly through the selective renderings by which finite judgment becomes possible.

Appendix A: Provisional Definitions

The following definitions are provisional but load-bearing. Their purpose is not lexical finality, but structural clarity. They are meant to stabilize the paper's working vocabulary while keeping distinct the principal objects, transitions, and conditions through which mediated judgment becomes possible.

Case

A *case* is the practical target of judgment: the person, situation, claim, event, process, or problem toward which action is directed. It is the paper's ordinary working term for what may be stated more formally as a world-fragment. A case is never identical to any one rendering of it.

World-Fragment

A *world-fragment* is the portion of reality toward which judgment is directed. It is the more formal counterpart to *case*. The term is useful because not every target of judgment is a simple object. A world-fragment may be a person, a process, a relational situation, an institutional condition, or a compound practical circumstance. In the main prose, *case* is usually preferable for readability, but the underlying point is the same: what judgment acts toward is not identical to the representation through which it becomes tractable.

Operative Representation

An *operative representation* is the rendering through which judgment becomes action-guiding. It may take the form of a file, category, score, profile, thresholded record, case summary, dashboard entry, model output, or other structured rendering stable enough for procedure, communication, or intervention.

A representation is *operative* when consequences for a case are materially routed through it within a judgment process. This criterion matters because not every representation is operative in the relevant sense. A rendering may exist, matter, or influence perception without yet functioning as the object over which consequences are organized. An operative representation is one that has entered the path of judgment in a way that helps determine what follows.

Access Trace

An *access trace* is a measurement, report, testimony, event record, perceptual uptake, document, or other trace through which some aspect of a case becomes available. The concept is meant to remain closer to access than to settled evidentiary force. A trace is something through which the case becomes available at all, not yet something that automatically counts in support of a judgment.

Evidentiary Interpretation

An *evidentiary interpretation* is an access trace taken under a standard of relevance, admissibility, credibility, or sufficiency. The point of this term is to mark that evidence is not simply found. A trace becomes evidence only within some regime that specifies what it is taken to show, how strongly it counts, and whether it counts at all.

The same trace may function as strong evidence in one setting, weak evidence in another, and irrelevant material in a third. The term therefore marks the transition from available trace to action-relevant support.

Conditioned Cut

A *conditioned cut* is the structured selection-regime through which a case is rendered into a usable representation under finite conditions. It determines what distinctions are preserved, what is backgrounded, what format the case must take, and what becomes tractable enough for judgment or action.

A cut is *conditioned* because it is shaped by constraints, tools, procedures, purposes, institutional formats, and practical stakes. The concept matters because it blocks the idea that simplification is merely an accidental defect. Selective rendering is part of what makes mediated judgment possible at all.

Judgment Regime

A *judgment regime* is the rule, procedure, institutional logic, decisional workflow, or evaluative process that acts over an operative representation. It specifies how the rendering is interpreted, what thresholds matter, what outputs are treated as decisive, what override rules exist, and how consequences attach.

The term is meant as a family concept rather than as one narrow mechanism. A judgment regime may be highly formal, as in a statutory test or algorithmic threshold, or more weakly formalized, as in a standardized workflow or patterned administrative practice.

Standing Condition

A *standing condition* is the condition under which a representation is allowed to enter judgment as action-relevant. It marks what gives a rendering enough recognized status to be treated as admissible, usable, or institutionally legible within a judgment regime.

Standing should be distinguished from decisional force. A representation may have standing — it may count, be visible, or be admitted — without yet being outcome-determining.

Decisional Force

Decisional force is the degree to which an operative representation actually governs what follows in a case. A representation has decisional force when outcomes are materially routed through it, absent override, exception, or correction.

This concept separates mere relevance from consequence-shaping power. Not everything that counts in a case governs it equally. Decisional force marks the difference between being part of the record and helping determine the result.

Scope Condition

A *scope condition* is the range within which an operative representation remains adequately action-guiding. Scope marks where the rendering can legitimately support judgment and where its authority begins to degrade.

This term matters because many failures are not failures of total falsehood. They are failures of travel beyond legitimate use. A representation may be useful, orderly, and even well-supported for one task while becoming misleading when exported to another task, scale, or consequence context.

Residue

Residue is the structured remainder relative to a rendering: what the operative representation omits, flattens, merges, backgrounds, or carries badly. Residue is not a mystical outside. It is the ordinary remainder produced whenever finite action depends on selective stabilization.

The concept blocks the recurrent confusion in which a useful rendering is treated as if it had therefore exhausted the case. What remains outside the rendering need not be unknowable in order to be consequential.

Constraint

Constraint refers to unavoidable limits on time, staffing, memory, money, computational capacity, coordination, training, attention, and interpretive bandwidth. Constraints help explain why simplification, standardization, and compression are unavoidable features of finite judgment.

The term should be used carefully. Constraints explain why systems cannot process cases in unconstrained fullness. They do not by themselves justify every cost imposed by the way those systems simplify.

Burden

Burden refers to the costs imposed on those who must live under, navigate, satisfy, or correct mediated judgment. Burdens may be documentary, procedural, interpretive, temporal, emotional, or logistical.

The distinction between constraint and burden is central. Constraints are conditions of finite action. Burdens are distributive outputs of how those conditions are handled. A system may face real limits while still shifting disproportionate reconciliation costs onto those with the least slack or power.

Traceability

Traceability is the degree to which a mediated judgment leaves inspectable traces of how a case was rendered, what counted, how the judgment regime operated, and where a result came from. Traceability does not require perfect transparency. It requires enough record of formation and action that challenge, review, and diagnosis remain possible.

The concept matters because correction depends on more than the abstract possibility of appeal. If the path from case to representation to consequence is opaque, correction becomes theatrical rather than usable.

Correction Path

A *correction path* is a pathway through which model–case mismatch can be surfaced and revised. This includes appeal, override, amendment, audit, rollback, recalibration, exception-handling, second review, and redesign.

It is useful to distinguish two levels of correction:

- **case-level correction:** revision, appeal, amendment, override, or exception in a particular case;

- **system-level correction:** audit, recalibration, retraining, redesign, or rule change aimed at patterned failure.

Correction matters because if mediated judgment is structurally fallible, revision cannot be treated as an optional kindness added after the fact. It is part of what keeps judgment answerable to the cases it governs.

Structural Fallibility

Structural fallibility is the standing possibility of patterned misfit arising from finite action through selective operative representations. It does not mean that judgment always fails, nor that all representations are equally unreliable. It means that mismatch is built into the condition of mediated action itself and therefore cannot be treated as a rare anomaly.

Structural fallibility may take the form of stale records, proxy drift, threshold distortion, omitted context, misclassification, forced commensuration, or cross-scale misfit. What unifies these phenomena is not one mechanism of error, but the fact that finite action proceeds through selective renderings rather than unconstrained contact with the case in full.

Reflexive Consequence

Reflexive consequence refers to the way mediated judgment can reshape the very case it governs. Once a rendering is tied to action, it may alter treatment, opportunities, institutional status, self-understanding, or future traces. The representation does not merely describe the case; it enters the case's ongoing organization.

This concept is extension-level rather than strictly required in every application. It becomes especially important where judgments feed back into the conditions they purport only to register.

Responsible Mediated Judgment

Responsible mediated judgment is mediated judgment that remains answerable to the structural conditions under which it operates. At minimum, this requires recognizing that action proceeds through operative representations rather than over the case in unconstrained fullness, that such representations are selective and scope-bounded, that burdens arise when mismatch is exported downward, and that traceability and correction remain available.

This is not yet a full ethical theory or final account of legitimacy. It is the paper's minimal practical standard for judgment that does not deny or conceal its own mediation.

Appendix B: Formal Sketch

This appendix provides a compact structural sketch of *Mediated Judgment Under Constraint*. It is not a complete formal calculus. Its role is narrower: to state the theory's core objects, relations, and derived pressures in a form more explicit than prose alone allows.

B.1 Core Objects

Let

$$W$$

denote a case or world-fragment.

Let

$$T(W) = \{t_1, t_2, \dots, t_n\}$$

denote the set of access traces available for W , where each t_i may be a measurement, report, testimony, event record, perceptual uptake, document, or other trace through which some aspect of W becomes available.

Let

$$E_J : T(W) \rightarrow \mathcal{E}$$

denote an evidentiary interpretation function relative to a judgment regime J . This function maps available traces into an evidentiary space \mathcal{E} , where traces are treated under standards of relevance, admissibility, credibility, or sufficiency.

Let

$$C$$

denote a conditioned cut. A cut is the selective rendering regime through which W becomes representable under finite conditions.

Let

$$R = C(W, T(W), E_J(T(W)))$$

denote the resulting representation of the case. R is the rendering of W actually made usable for judgment.

Let

$$\Delta = \text{Res}(W \mid C)$$

denote the residue relative to the cut C , that is, the structured remainder not preserved, flattened, merged, backgrounded, or badly carried by R .

Let

$$J$$

denote a judgment regime, that is, the rule, procedure, workflow, or institutional logic acting over

R .

Let

$$S(R, U) \in \{0, 1\}$$

denote the scope condition of R for use-context U , where $S(R, U) = 1$ means that R remains adequately action-guiding in U , and $S(R, U) = 0$ means that it does not.

Let

$$\Sigma(R, J) \in \{0, 1\}$$

denote the standing condition of R under J , where $\Sigma(R, J) = 1$ means that R is admitted or recognized as action-relevant within the judgment regime.

Let

$$\Phi(R, J) \in [0, 1]$$

denote the decisional force of R under J , where higher values indicate greater outcome-shaping force.

Let

$$A$$

denote the action or decision produced under mediated judgment.

Let

$$B$$

denote the burden distribution generated by the mediated judgment process.

Let

$$K$$

denote the correction path, that is, the set of possible review, appeal, override, amendment, audit, recalibration, rollback, or redesign operations available relative to A , R , and W .

B.2 Formation of an Operative Representation

The basic formation sequence is:

$$W \longrightarrow T(W) \xrightarrow{E_J} \mathcal{E} \xrightarrow{C} R$$

This states that a case does not become action-guiding all at once. It first becomes available through traces, then those traces are interpreted as evidence under a judgment regime, and only then rendered into a representation usable for action.

A representation is *operative* when it is not merely present, but materially enters the consequence path of a case. Formally:

$$\text{Opr}(R, J) \iff \Sigma(R, J) = 1 \wedge \Phi(R, J) > 0$$

That is, a representation is operative if and only if:

1. it has standing within the judgment regime,
2. and it has non-zero decisional force.

This blocks the confusion between a representation that merely exists in the background and one through which consequences are actually routed.

B.3 Mediated Action

The central claim of the paper can now be stated schematically:

$$A = J(R) \quad \text{where } R = C(W, T(W), E_J(T(W)))$$

This means that action is not taken on W in unconstrained fullness, but on the representation R formed through selective rendering.

A more explicit version is:

$$A = J(R \mid \Sigma, \Phi, S)$$

meaning that the judgment regime acts on R under conditions of standing, decisional force, and scope.

B.4 Scope and Overextension

A representation may function adequately in one use-context and fail in another. Let U denote a use-context. Then

$$S(R, U) = 1$$

means that R is adequately action-guiding in U , while

$$S(R, U) = 0$$

means that it is not.

We may define overextension as:

$$\text{Overext}(R, U) \iff S(R, U) = 0 \wedge \Phi(R, J) > 0$$

That is, overextension occurs when a representation continues to exercise decisional force in a use-context outside the conditions under which it remains adequately action-guiding.

This captures one of the paper's key points: many failures are not failures of total falsehood but failures of illegitimate travel.

B.5 Structural Fallibility

Because R is produced through a selective cut C , residue is always possible:

$$\Delta = \text{Res}(W \mid C)$$

Structural fallibility can therefore be stated as:

$$\forall R \exists \Delta \text{ such that } R \neq W \text{ and } \Delta \text{ may become action-relevant}$$

This should not be read as a claim that all judgments are equally bad or equally unreliable. Its point is narrower: where action depends on selective rendering, patterned misfit is a standing possibility of the structure itself.

A more diagnostic form is:

$$\text{SF}(R, J, U) \iff \text{Opr}(R, J) \wedge \exists \Delta \text{ such that } \Delta \text{ is relevant in } U$$

So structural fallibility holds where:

1. the representation is operative,
2. and some residue becomes relevant to the context in which action is taken.

B.6 Burden Distribution

Mismatch does not vanish once judgment proceeds. It is absorbed somewhere. Let

$$B = \text{Dist}(A, R, \Delta)$$

denote the burden distribution generated by acting through R despite Δ .

The key structural claim is:

$$\text{SF}(R, J, U) \Rightarrow B \neq 0$$

under ordinary non-trivial conditions.

In plain terms, when an operative representation misfits the case in a consequential way, reconciliation costs do not disappear. They are distributed as burdens.

These burdens may take several forms:

$$B = \langle B_d, B_p, B_i, B_t, B_e, B_l \rangle$$

where:

- B_d = documentary burden,

- B_p = procedural burden,
- B_i = interpretive burden,
- B_t = temporal burden,
- B_e = emotional burden,
- B_l = logistical burden.

This decomposition is illustrative rather than exhaustive, but it marks the principal channels through which mediated judgment becomes lived and social rather than merely informational.

B.7 Correction

If structural fallibility is built into mediated judgment, then correction must also be structurally available.

Let

$$K = \text{Corr}(A, R, \Delta, B, J)$$

denote the set of correction paths available relative to an action A , operative representation R , residue Δ , burden B , and judgment regime J .

We may distinguish:

$$K = K_c \cup K_s$$

where:

- K_c = case-level correction paths,
- K_s = system-level correction paths.

Case-level correction includes

$$K_c = \{\text{appeal, override, amendment, exception, second review}\}$$

System-level correction includes

$$K_s = \{\text{audit, recalibration, redesign, retraining, rule revision}\}$$

The theory's structural claim is:

$$\text{SF}(R, J, U) \Rightarrow K \neq \emptyset$$

for responsible mediated judgment.

This is not yet a full ethical claim. It is the paper's architectural point: if mediated action is structurally fallible, then a system with no effective correction path is built as though its operative renderings were self-exhausting.

B.8 Traceability

Correction depends on traceability. Let

$$\text{Tr}(W, R, J, A) \in [0, 1]$$

denote the degree of traceability in the mediated judgment process.

Higher traceability means that the path from case to traces to representation to action remains sufficiently inspectable for challenge and revision.

A minimal dependency claim is:

$$K \text{ usable} \Rightarrow \text{Tr}(W, R, J, A) > 0$$

That is, correction cannot be more than theatrical if the route by which judgment was formed is wholly opaque.

B.9 Reflexive Consequence

In some domains, mediated action reshapes the very case it governs. Let

$$W' = F(W, A, R)$$

denote the transformed case after mediated action feeds back into it.

This yields the reflexive sequence:

$$W \xrightarrow{C} R + \Delta \longrightarrow A = J(R) \longrightarrow W' = F(W, A, R)$$

The point of this extension is that mediated judgment may not merely register cases but alter their future structure, future traces, and future governability.

B.10 Compact Theory Schema

The full structural sketch may now be compressed into:

$$W \xrightarrow{T, E_J, C} R + \Delta$$

$$\text{Opr}(R, J) \iff \Sigma(R, J) = 1 \wedge \Phi(R, J) > 0$$

$$A = J(R \mid \Sigma, \Phi, S)$$

$$\text{Overext}(R, U) \iff S(R, U) = 0 \wedge \Phi(R, J) > 0$$

$$\text{SF}(R, J, U) \iff \text{Opr}(R, J) \wedge \exists \Delta \text{ relevant in } U$$

$$\begin{aligned} \text{SF}(R, J, U) &\Rightarrow B \neq 0 \\ \text{SF}(R, J, U) &\Rightarrow K \neq \emptyset \\ K \text{ usable} &\Rightarrow \text{Tr}(W, R, J, A) > 0 \\ W' &= F(W, A, R) \end{aligned}$$

B.11 Interpretive Summary

The formal point of the sketch is simple.

A case becomes actionable only through selective rendering. Once that rendering becomes operative within a judgment regime, it may acquire standing and decisional force. Because the rendering is selective, it leaves residue. Once consequences attach, that residue may become action-relevant. When it does, mismatch is expressed not only as epistemic defect but as distributed burden. If the system is to remain answerable to the case it governs, it must preserve traceability and correction.

That is the structural core of *Mediated Judgment Under Constraint*.

References

- [1] Geoffrey C. Bowker and Susan Leigh Star. *Sorting Things Out: Classification and Its Consequences*. MIT Press, Cambridge, MA, 1999.
- [2] Virginia Eubanks. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Press, New York, 2018.
- [3] Miranda Fricker. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press, Oxford, 2007.
- [4] Ian Hacking. The looping effects of human kinds. In Dan Sperber, David Premack, and Ann James Premack, editors, *Causal Cognition: A Multidisciplinary Debate*, pages 351–383. Clarendon Press, Oxford, 1995.
- [5] Donna Haraway. Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist Studies*, 14(3):575–599, 1988.
- [6] Donald P. Moynihan, Pamela Herd, and Hope Harvey. Administrative burden: Learning, psychological, and compliance costs in citizen-state interactions. *Journal of Public Administration Research and Theory*, 25(1):43–69, 2015.
- [7] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, 2019.
- [8] Cathy O'Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown, New York, 2016.

[9] Michael Polanyi. *The Tacit Dimension*. Routledge & Kegan Paul, London, 1966.

Rights and Contact

Copyright © 2026 David T. Swanson.

This work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0).

DOI: [10.5281/zenodo.19078805](https://doi.org/10.5281/zenodo.19078805)

For questions, comments, or citation inquiries, contact: dswanson903@gmail.com