

Normative Information Preservation in Representational Renormalization: An Investigation of Contestability and Remedial Authority

Wanhong HUANG
huangwanhong@serendip.ngo

Abstract

Institutional representations condense heterogeneous histories into classifications, records, and indicators that may shape consequential decisions. This paper asks when an institution can responsibly rely on a reduced description while retaining the capacity to recognize, examine, and respond to challenges that its own categories may obscure. The central thesis is conditional: representational adequacy for a declared question, evidential access, review competence, effective correction, and the legitimacy of the governing criteria require distinct warrants. None alone establishes the others. Elementary results on factorization, approximation, indistinguishability, and guarded reachability identify where particular inferences fail under stipulated assumptions; they establish neither moral verdicts nor institutional powers. Renormalization supplies a bounded analogy rather than a demonstrated physical scaling law for institutions. Drawing on work in classification, epistemic injustice, sociotechnical abstraction, and accountability, the paper proposes a question-indexed, seven-part cross-stage inquiry covering recognition, evidence, access, competence, transfer, response, and frame review. Its incremental value lies in locating the precise stage and missing warrant between an identifiable representational distinction and a potentially consequential institutional response, rather than introducing new mathematics or generic audit principles. Fourteen propositions are editorially classified by mathematical role and inferential limits. Constructed cases consider excluded claims, feedback, protected evidence, referrals, competing remedies, and challenges to evaluative rules. A prospective protocol compares the proposed integrated register with competent existing procedures, including the possibility of equivalent performance, greater burden, unsafe disclosure, and inconclusive effects. No empirical implementation, legal validation, causal estimate, or independent moral proof is claimed. The result is a bounded investigative method whose usefulness depends on independently warranted questions, reliable evidence, legitimate institutional authority, and comparative performance.

Keywords: representational renormalization; normative information preservation; contestability; remedial authority; coarse-graining; epistemic justice

Research and Publication Notice

Status. Final editorial author-review candidate, version 3.1, 24 October 2026. The nine-section manuscript and its constructed replication appendix, unimplemented comparative-evaluation protocol, and cross-paper concordance appendix have been prepared as an author-review candidate. The author has selected the present title and subtitle; approval for public dissemination remains a separate decision. The document is neither an independently peer-reviewed article nor a validated empirical study; it has not been designated by the author as a public release.

Theoretical scope. The investigation concerns representational decisions and their institutional and epistemic consequences. Ordinary summaries, statistical aggregation, discretization, and physical renormalization have related but distinct mathematical meanings. The expression “representational renormalization” identifies the paper’s bounded multiscale analogy; no technical renormalization flow for normative systems is established.

Ontological and temporal restraint. Relational histories and scale mappings are analytical representations. Persons and institutions require historically interpreted relational structure and are not stipulated as primitive vertices. Event precedence, record order, narrative order, and model parameter order require independent specification.

Citation scope. The ledger distinguishes publisher abstracts, accessible primary passages, and institutional or standards documents. It records bounded source-to-claim checks rather than exhaustive passage-by-passage verification. Bibliographic presence alone is not a guarantee of substantive support.

Human responsibility and licence. OpenAI’s ChatGPT assisted with drafting, revision and technical validation. The author retains responsibility for correctness, verification, approval and dissemination. Copyright © 2026 Wanhong Huang. The author intends to offer the original text under CC BY-NC 4.0 following author approval, subject to third-party rights. <https://creativecommons.org/licenses/by-nc/4.0/>.

1 Introduction

Institutional representations make complex records tractable for inquiry and collective decisions. Researchers, administrators, designers, and public authorities select observations, construct variables, define classifications, and publish summaries for particular purposes. Each operation establishes boundaries around the information available for subsequent interpretation. In consequential settings, these boundaries may affect whether a contribution is recognized, a claim becomes reviewable, or a burden is considered in an institutional decision. The research object of this paper is the responsibility associated with constructing, authorizing, and using those representations.

Three constructed examples illustrate the problem. A service programme reports its rate of completed provision but does not retain a field for unsuccessful appeals. A research infrastructure credits a consortium while its contribution categories leave some forms of labor unrecognized. An environmental model reports an aggregate burden while omitting distinctions needed to examine geographically uneven or irreversible effects. These examples concern different kinds of potential omission. They invite inquiry into the justification of the chosen informational purpose, the histories admitted to the record, and the authority to revise a representation when consequential claims become visible. They are analytical constructions, not accounts of observed institutions.

Existing scholarship establishes substantial foundations for this inquiry. Bowker and Star analyze how categories and standards organize social and institutional visibility [7]. Fricker examines wrongs associated with credibility and interpretive resources in the participation of persons as knowers [16]. Selbst and colleagues identify difficulties created by an overly restrictive sociotechnical abstraction when technical assessments are applied to institutional settings [43]. These contributions already make classification, epistemic standing, and the social consequences of abstraction legitimate objects of investigation. The present paper proposes a narrower analytical interface: the identification of a consequential distinction lost at a particular representational stage, the evidence potentially available to recover it, and the competent process through which a justified correction may be considered.

The inquiry also engages the level dependence of normative description. List and Valentini argue that political normative facts may be coarser-grained than moral facts [29]. Their argument concerns the relationships among normative descriptions and appropriate levels of political analysis. The present investigation concerns the accountable practices by which institutions select, maintain, and deploy particular descriptions. It follows Paper I's investigation of generative diminishment and asymmetric relations, Paper II's analysis of credible commitments and remedies, and Paper III's study of preserving normative claims across institutional scales. Paper IV examines the institutional production and consequential circulation of the descriptions on which such inquiries may depend. These are related research objects with distinct questions of justification.

The inherited expression *renormalization* requires terminological restraint. Renormalization-group analysis in physics involves specified transformations, parameter behavior, and, in appropriate settings, fixed points and critical phenomena [48]. Statistical aggregation, the reduction of historical records to institutional categories, and the construction of effective descriptions require their own definitions. The manuscript uses the technical physics literature as a carefully bounded comparison. It does not derive a renormalization-group flow for ethical systems, nor assume that an institutional summary has the properties of a physical coarse-graining transformation. The working title remains provisional for this reason.

The thesis has two complementary parts. First, a reduced description is *question-adequate* only relative to a specified domain, admissible interpretation, and task; exact recovery of one answer says nothing by itself about an omitted question or the institutional standing of its proponent. Second, when a representation is used consequentially, an adequate descriptive channel does not imply an adequate route for challenge and action. Access to a record, authority to reconsider a decision, and a substantively justified response are different claims. Conversely, failure of public recovery does not prove an institutional wrong if reliable, lawfully accessible supplementary evidence and competent review provide an appropriate route. The proposed method organizes the evidence needed to distinguish these possibilities.

Its normative adequacy is deliberately not defined as a numerical function of information retained or procedures completed.

The paper's contribution can be evaluated through a discriminating comparison: given the same institutional episode and evidential access, does a question-indexed examination across admission, reduction, protected inspection, authorized review, downstream reliance, and rule revision identify a materially important and properly actionable omission that established case review or provenance practice would miss? If existing procedures already reveal and resolve the relevant issue at equal or lower cost and risk, the proposed combined register has not demonstrated added practical value. This comparison makes the paper's institutional proposal contestable while leaving ethical interpretation and legal duties to the appropriate arguments and authorities.

The proposed accountability interface spans distinguishable decisions. A history-admission rule determines which claims enter an evidence domain; a variable-selection rule specifies which attributes are retained; a reduction can make previously distinguishable histories observationally equivalent; a disclosure policy determines which audiences may inspect the retained distinctions. Consequential users may then rely on those representations to make decisions, while affected participants may challenge their construction or use. Correction at one stage may be immaterial to a downstream decision unless reliable information reaches an authorized reviewer. These stages connect information loss to institutional conduct while leaving their legal and ethical evaluation open to further argument.

A relational history here is an analytically specified record of events, couplings, observations, institutional practices, and interpretations. Its representation presupposes neither a universal linear ordering of time nor an ontologically privileged present. Where cellular or two-complex illustrations are relevant, vertices denote events and edges couplings; the interpretation of faces is stipulated locally. Subjects and institutions are historically interpreted configurations rather than primitive vertices. This formal vocabulary organizes questions about evidence without supplying an ontological proof of social or moral reality.

The mathematical analysis addresses the preservation of answers relative to a declared domain and family of questions. When two admitted histories receive the same reduced description but different answers to a relevant question, the reduced description alone is insufficient for reconstructing that answer. The conclusion depends on the histories and interpretive questions that have been admitted. Exclusion may therefore occur before any reduction, through testimonial access, classification, or the language available to formulate a claim. Further evidence may reopen an inquiry even where an earlier representation has lost a distinction. The validity of such evidence and the authority to use it require independent examination.

More information is not invariably a more responsible public representation. Confidentiality, risk of retaliation, evidential quality, procedural fairness, and context-dependent rules of disclosure may constrain admissible refinements. Nissenbaum's contextual-integrity account offers a basis for assessing information flows in their social settings [33]. Related work by Jasanoff and Saltelli and colleagues emphasizes uncertainty, framing, participation, and the responsible use of models in consequential public decisions [23, 42]. Their contributions provide significant precedents; the present analysis does not claim novelty for transparency, participation, documentation, or auditing in themselves.

The manuscript advances a conceptual and exploratory formal proposal. It distinguishes descriptive adequacy, evidential support, causal explanation, institutional competence, and substantive ethical justification. A proposed cross-stage assessment traces question-relevant information from admission and classification through reduction and consequential use, while identifying possible channels for challenge and correction. Its elementary formal results are applications of established mathematics. Its potential methodological contribution lies in organizing the evidential and institutional interfaces among these stages; that contribution requires comparison with established qualitative inquiry, legal procedures, and organizational audit methods. The hypothetical cases are designed for analytical discrimination and provide no empirical validation of effectiveness.

The synthesis in Section 9.19 makes the paper's architecture explicit. A declared question, its evidential preservation, protected reviewer access, authorized competence, effective transmission, permissible

disposition and the governance of evaluation criteria are independently assessable elements of an institutional claim pathway. Their distinction explains why accurate records may coexist with inaccessible review, why a valid correction may fail to reach a consequential decision, and why complete internal case logs may leave rule-level exclusions undetected. The proposed integrated register organizes these questions without reducing them to a single ethics metric or transferring authority from competent institutions to a formal model.

The analysis proceeds through scale selection and representational authority in Section 2, admission and variable selection in Section 3, and preservation conditions in Section 4. Subsequent sections investigate historical annotations, epistemic exclusion, contestation and corrective propagation, and the accountability conditions for consequential use. Section 9 then examines conflicts, evidence requirements, intellectual precedents, and conditions under which the proposed methodology would add little or no value. The resulting inquiry treats representational governance as historically situated and revisable, with attention to heterogeneous claims and the generative conditions of human and nonhuman life.

2 Scale Selection and Representational Authority

This section establishes the analytical objects required to evaluate representational choices. Its objective is to distinguish different kinds of reduction, specify the decisions through which a descriptive scale acquires practical authority, and delimit the informational claims that can be supported by elementary formal results. The method combines conceptual distinctions with a partial formalization of admitted histories and a constructed institutional example. The resulting vocabulary identifies questions for later normative and institutional inquiry; it leaves their ultimate justification open.

2.1 Aggregation, coarse-graining, and effective description

Different operations often described as simplification have different mathematical and practical roles. *Aggregation* combines observations through an operation, such as a count, sum, or average. *Coarse-graining* groups distinguishable fine descriptions into equivalence classes or maps them to a coarser representational space. *Discretization* replaces aspects of a potentially continuous model with a selected finite or countable representation. An *effective description* is a model whose variables and transformation rules are chosen for explanatory or predictive use at a particular level of inquiry. A *renormalization-group construction*, in a technical physical setting, requires an appropriately defined family or transformation structure with relevant parameter behavior and consistency relations [48].

These distinctions concern the structure and use of descriptions. They do not rank the methods by moral worth or epistemic reliability. Averaging can be useful for estimating a population-level indicator. A coarse-grained category can simplify an explanation while retaining precisely the distinctions needed for an institutional task. A fine-grained record can retain irrelevant or sensitive detail whose collection and disclosure require justification. The consequential issue is the relation among the descriptive method, its purpose, the background assumptions, and the institution that makes it operational.

Table 1 compares the kinds of claims licensed by different forms of reduction. The final column identifies the additional information needed when the method enters public or institutional decision-making.

The terms in Table 1 are also context-sensitive. The statistical aggregation of records may itself form one stage of a coarse-graining operation. An effective description may employ aggregation, discretization, or both. The distinctions specify different claims that an author should establish rather than disjoint families of activity. Consequently, one modeling pipeline can create several distinct opportunities for institutional scrutiny.

Table 1: Descriptive transformations and their respective accountability questions.

Operation	Technical question	Accountability question
Aggregation	Which statistic combines the admitted observations?	Whose observations enter the population and whose burdens are hidden by the statistic?
Coarse-graining	Which fine descriptions share a reduced state?	Which claims become indistinguishable, and who can request refinement?
Discretization	Which resolution and boundary conventions are imposed?	Which changes and threshold crossings may be missed?
Effective modeling	Which variables and transformations answer the declared task?	What warrants the declared task, domain, and deployment conditions?
Technical renormalization	Which transformations define scale flow and parameter behavior?	What interpretation and evidential warrant support transporting the formalism outside its established domain?

2.2 Representational policies and admitted histories

A formal object becomes useful once its elements correspond to identifiable descriptive decisions. Let $\mathcal{H}$ denote a candidate space of enriched histories available to an investigator under a declared interpretive scheme. A history may include records of events, couplings, institutional classifications, relevant testimony, and admissible interpretations of subjects or obligations. The domain $D \subseteq \mathcal{H}$ represents the histories admitted for a particular analytical operation. Admission may depend on evidence access, institutional jurisdiction, consent, measurement technologies, or a modeling convention. Thus D is a substantive component of the represented inquiry, rather than an unproblematic list of everything that exists.

Let $\mathcal{T}$ be a coarser description space and let $U : D \rightarrow \mathcal{T}$ be a specified reduction. A normative question q is associated with an answer space A_q and an interpreted evaluation $J_q : D \rightarrow A_q$. The answer space might distinguish eligibility to appeal, obligations of an authorized institution, or another question with stated assumptions. We initially take these evaluations as deterministic total functions to isolate a modest factorization result. Partial, disputed, and set-valued answers require further analysis.

For institutional purposes, it is helpful to regard the map U as one component of a broader *representational policy*. A minimal policy record is stated in Equation 1.

$$\pi = (D, \iota, U, \mathcal{Q}, \mathbf{P}, \mathbf{R}). \quad (1)$$

Here ι denotes the interpretive and evidential conventions by which records and subjects are described; $\mathcal{Q}$ denotes a declared family of questions; $\mathbf{P}$ records provenance and institutional authorization; and $\mathbf{R}$ identifies available procedures for scrutiny or revision. The tuple is a bookkeeping device. It makes a set of decisions explicit without assuming that these decisions are exhaustive or that their mere documentation justifies their effects. Equation 1 should be read as a candidate minimum audit interface, subject to modification according to the institutional setting.

Each component of π has a different possible ethical significance. The admitted domain affects the population to which an analysis applies. The interpretive conventions influence which accounts can become recognizable evidence. The reduction specifies which distinctions remain available downstream. The question family partly determines what counts as adequacy. Provenance permits inquiries into where assumptions originated and by whom they were authorized. Revision procedures affect whether omitted or misrepresented subjects have a realistic opportunity to contest the representation. None of these roles entails that a single institution or individual should control them all.

The interpretation of subjects is particularly important in relational models. An event-level cellular complex may supply vertices interpreted as events, edges interpreted as couplings, and faces given locally declared primitive semantics. A person, community, or institutional body is then an interpreted historical structure rather than a primitive vertex. The availability of a vertex does not, by itself, establish an

institutional subject with procedural standing. A modeler who equates a selected event category with a complete claimant population has made an additional interpretive decision; its warrant belongs in ι and may need testing against excluded or disputed accounts.

A related complication concerns ordering. A data set may be arranged by entry date, a causal graph may record a justified partial precedence relation, and a narrative may impose a particular sequence on disputed testimony. These structures have different meanings. An aggregation over dated observations cannot by itself resolve contested causal ordering or the historical priority of an obligation. Such claims depend on specific evidence and interpretive rules. The present framework therefore treats event precedence, record order, narrative sequence, and modeling steps as independent relations whenever they become material to the argument.

2.3 Query-relative preservation and its limited warrant

The preservation of a declared question can be stated precisely within the stipulated function model. A reduced evaluation $\bar{J}_q$ exists on the attained coarse descriptions $U(D)$ when the fine evaluation is constant within each fiber of U . The elementary criterion is given in Equation 2.

$$U(h) = U(h') \implies J_q(h) = J_q(h') \quad \text{for } h, h' \in D. \quad (2)$$

Equation 2 is evaluated for a particular question q and a particular admitted domain D . It identifies whether the answer to that question can be recovered from the reduced description alone. When the condition holds, the reduced evaluation is defined by setting $\bar{J}_q(U(h)) = J_q(h)$; the condition ensures that the definition is independent of the representative history selected from the fiber. Conversely, the existence of such an evaluation implies the condition. This is the elementary factorization criterion already developed in the author's prior exploratory two-complex work and employed in Paper III. It is restated here solely to support the ethical analysis of representational policy.

The fact that Equation 2 holds supplies a formal answer about functional sufficiency. It does not establish that J_q is justified, that the admitted histories faithfully represent affected subjects, or that an institution is authorized to act using $\bar{J}_q$. The same answer may be preserved for a question framed by an unfair eligibility rule. A wider domain may contain recognized claims that were absent from D . Even within D , an available question family may be too narrow to address a relevant ecological, relational, or procedural burden. These concerns call for a separate examination of the authority to select Q and the evidential conditions governing D and ι .

The finite diagnostic in Equation 3 makes an observed query-specific loss explicit without assigning a scalar ethical value to it.

$$\mathcal{W}_q(\pi) = \{(h, h') \in D^2 : U(h) = U(h'), J_q(h) \neq J_q(h')\}. \quad (3)$$

Each pair in $\mathcal{W}_q(\pi)$ is a witness to the failure of exact answer preservation for the stated question. Equation 3 follows directly from Equation 2. A nonempty witness set is an informational diagnosis, while its institutional significance depends on the content of the answer difference, the consequences of the reduction, and the available methods of review. For infinite or incompletely sampled historical spaces, an empty observed witness set may reflect limited access rather than complete preservation.

The diagnostic also reveals a tension within normative model design. If the declared question family is enlarged, a reduction previously adequate for the narrower family may fail. If the question family is restricted, the same map may appear adequate without any change in the historical records. A policy of selecting only easily preserved questions would therefore risk turning representational convenience into a surrogate for normative relevance. The responsibility to justify the question family is conceptually independent of the responsibility to calculate preservation correctly.

2.4 An institutional case of representational substitution

Consider a hypothetical environmental restoration programme operating through a federation of local councils. Its published summary reports the number of completed interventions and a combined index of

Table 2: Questions arising from the constructed restoration-programme comparison.

Analytical domain	Focused question	Evidence or argument required
Formal preservation	Do the reduced records determine eligibility for review?	Admitted histories, map definition, answer conventions, and witnessed distinct answers
Institutional process	Who selected eligibility, reviewed disputes, and authorized reliance on the summary?	Applicable competences, decision records, consultation, appeals, and implementation histories
Ethical justification	Are the affected claims appropriately recognized and is the use of the reduction defensible?	Substantive normative principles, relevant consequences, counterarguments, and affected-subject interpretations

ecological recovery. For an authorized aggregate performance report, these values may be informative. Suppose that two possible programme histories yield the same published totals. In the first, a small community successfully obtained a review of a restoration decision; in the second, members of that community were excluded from the procedure through an eligibility classification. The situations are stipulated, and the numerical equality is fixed by construction.

Let h_a and h_b denote the two constructed histories. The aggregate map satisfies $U(h_a) = U(h_b)$, while the question q_{review} concerning effective eligibility to seek review has different answers. Equation 3 therefore identifies (h_a, h_b) as a loss witness. The same map may preserve a question concerning total recorded interventions. This comparison makes two observations available. First, a useful aggregate can be insufficient for a different institutional purpose. Second, deciding which of those purposes receives priority is a policy decision whose legitimacy cannot be inferred from an aggregate model’s arithmetic correctness.

A further distinction matters for the governance of that decision. Imagine that the model developers are employed by the federation, the eligible claimant population is specified in a separate institutional rule, and the public agency using the summary has authority to allocate restoration funds. The developer may choose the statistic within a technical remit; the federation may define data eligibility; the agency may determine the consequences of the resulting classification. The authority to represent, to determine standing, and to act on model output is distributed across these institutions. Ethical accountability therefore requires investigation of their linked responsibilities, rather than treating a single developer as the entire causal source of the outcome.

An adequate audit would ask who established the eligibility classification, whether affected communities could recognize and correct its effects, whether the aggregate report was represented as suitable for procedural review, and what evidence connects its use with consequential decisions. Competing explanations remain possible. A denied appeal might originate in a distinct legal requirement rather than the representational map. A resource shortage might explain an implementation failure irrespective of the summary’s granularity. A revised data category might improve visibility without altering the institution’s decision-making authority. These alternatives need to be considered before a researcher attributes an injustice to representational reduction.

Table 2 distinguishes three classes of questions prompted by the example. Their separation limits the tendency to infer an ethical conclusion from a formal diagnostic alone.

The constructed case also demonstrates why coarse-graining should be analyzed as a sequence of decisions rather than as a single abstract mapping. Domain admission, variable selection, model training or calculation, publication, circulation through institutions, and eventual deployment may occur under different permissions and forms of oversight. Provenance frameworks such as the W3C’s PROV-O provide a vocabulary for representing and exchanging information about entities, activities, and agents in such processes [47]. Provenance can assist an inquiry into how a summary came to exist. It cannot, by itself, certify the moral adequacy of the choices recorded or determine whether a disputed account received appropriate consideration.

2.5 Selection, authorization, and the distribution of responsibility

The authority to select a scale may be epistemic, administrative, legal, contractual, or practical. A research team might choose a resolution because it permits stable estimation from available measurements. An administrator may rely on a reporting category required by a public rule. A consortium may negotiate an aggregation procedure as part of an agreement. A public authority may deploy a model despite limited opportunities for challenge. Their responsibilities depend on the functions undertaken, what was reasonably knowable, the available alternatives, and the foreseeable consequences of the choice.

Distinguishing these forms of authority avoids an automatic equation of technical expertise with legitimate governance. An investigator can have substantial knowledge of statistical error while lacking authority to determine institutional standing. Conversely, a public body may have recognized legal competence to adopt a reporting convention while relying on evidence inadequate for its declared purpose. The institution's reason for adopting a representation should therefore be recorded and assessed in connection with its actual remit and the affected forms of participation.

The relevant ethical inquiry may concern duties of disclosure, consultation, correction, restraint, independent review, or documentation of uncertainty. The precise duties vary by context. A model used to explore a theoretical mechanism within a laboratory places different burdens on its authors from a model used to restrict public services or determine irreversible interventions. A prospective account of responsibility should examine the seriousness and reversibility of possible consequences, the distribution of expertise and power, feasible means of challenge, and the degree to which the model displaces alternative forms of evidence.

The selection process also has an informational political economy. Measurement and documentation require labor, maintenance, infrastructure, and interpretive conventions. Those who provide detailed observations may have limited authority over their conversion into institutional indicators. Eglash and collaborators' work on generative justice offers a pertinent ethical context by examining the conditions under which value-generating communities can retain greater control over production and the circulation of generated value [12]. The argument here concerns an analogous but separate issue: participation in defining the representational interfaces through which contributions and burdens acquire public or administrative visibility. The analogy requires context-specific evidence and entails no identity between data governance and Eglash and colleagues' proposed economic arrangements.

2.6 Privacy, disclosure, and protected opacity

Some lost information may deserve preservation for audit purposes while requiring protection from public release. A finer institutional record can disclose sensitive relationships, expose persons to retaliation, or create new surveillance opportunities. Ethical scrutiny of coarse-graining must therefore also assess the consequences of increasing descriptive granularity. A categorical preference for maximal information retention would neglect context-sensitive interests in privacy and safety.

Nissenbaum's account of contextual integrity emphasizes the norms governing information flows within different social contexts [33]. This insight supports treating preservation, access, and disclosure as distinct decisions. A confidential appeal record could remain available to an authorized reviewer while an aggregate public report reveals only a general count. Such a design may protect participants and retain the institutional possibility of challenge. Yet it also raises questions about the reviewer's independence, the reliability of access procedures, the availability of redress for wrongful withholding, and safeguards against selective disclosure. A policy may be formally traceable while practically inaccessible to the people who need its protections.

The privacy problem illustrates a general methodological limit. A query-preserving refinement of a reduced state may require additional distinctions, but an admissible refinement for institutional purposes depends on permissions, security, community expectations, and the legitimate scope of an inquiry. Adequate representation is therefore conditional on the relation between questions and the governance of information. Later sections will examine the balance between contestability and protected opacity, including circumstances where the relevant account can be represented through confidential witnesses,

protected audit procedures, or appropriately abstract descriptions.

2.7 Limits of the initial framework

The initial policy tuple and factorization diagnostic are deliberately restricted. They abstract from disputed interpretations, incomplete testimony, stochastic outputs, uncertain causal relations, and changes in the admitted domain. They do not define a theory of ethical aggregation or a unique allocation of institutional responsibility. Their intended function is diagnostic: to identify where a reduction is demonstrably unable to answer a declared question and to distinguish that result from questions of authority and legitimacy.

The framework will also need to account for historically evolving interpretations. What counts as an affected subject, an intelligible harm, or a permissible form of evidence can change through deliberation, conflict, institutional reform, and new observations. A model that was sufficient relative to one specified set of questions may require re-examination under another defensible interpretation. In the present framework, such changes are described through alterations of the domain, semantic conventions, questions, or procedures, with no assumption that the sequence is governed by a universal clock or privileged ontological present.

The analysis in Section 2 therefore supports a limited but consequential conclusion. Representational reduction is assessable through distinct formal, evidential, institutional, and ethical inquiries. Technical preservation results can reveal specific losses; they require additional arguments to establish appropriate institutional use. Section 3 examines variable and domain selection, including the conditions under which affected people can contest the definitions used to make their claims legible.

3 Selection of Variables and Admissible Histories

This section examines the production of the descriptive domain on which representational reductions operate. Its objective is to distinguish admission of a history, extraction of a variable, and authorization of a question, since each decision may affect the recognition of institutional claims. The analysis combines an explicit observation model, an elementary non-identifiability result, and a constructed environmental-governance comparison. It subsequently considers selection warrants, protections for sensitive evidence, rival explanations, and conditions under which the proposed diagnostic adds little to ordinary institutional inquiry. The formal construction identifies possible losses; its interpretation requires independent evidential and ethical argument.

3.1 Admission rules and the represented population

The admitted domain $D \subseteq \mathcal{H}$ introduced in Section 2 depends on substantive decisions about what counts as an eligible record or history. A domain may be constrained by geographic jurisdiction, legal standing, available measurement technologies, access to archives, consent conditions, or the practical resources of an inquiry. Such constraints have different justifications and consequences. Some limits make a research question tractable. Others protect vulnerable subjects from disclosure. Still others reflect historically established rules whose adequacy is contested. A transparent model should identify the particular function of each limit rather than treat all exclusions as interchangeable omissions.

An institution may report a population of *registered claimants* while a research question concerns the wider population *affected by a decision*. The difference is material even if the registration statistics are accurate. Eligibility to register may itself require documentation, language competence, residence at an administrative address, or recognition under a category whose interpretation is disputed. As a result, records of persons who successfully entered a process may support strong claims about the process's registered population and weak claims about those unable to enter it. The World Bank's account of representation in data systems specifically identifies the absence of formal addresses, identification, or digital access among possible sources of invisibility in data collection [49]. The broader analytical point

is that the sampling frame, administrative population, and population relevant to a normative question require separate specification.

Classification has historically operated through practical infrastructures and standards, rather than through labels detached from institutional practices. Bowker and Star examine consequences of classification and standardization for visibility and social organization [7]. The present proposal retains that insight while asking a narrower methodological question: which differences in admitted histories can an authorized representation legitimately claim to have observed? Answering it requires a record of admission criteria, the grounds for exclusions, and the actual routes through which affected parties can challenge their categorization.

The significance of a domain boundary also depends on the purpose of the representation. A study reporting the number of registered appeals within a specified period may appropriately use the registry as its observational domain. A claim that the same registry measures the availability of appeal to all affected communities requires a different evidential warrant. The first inquiry asks about a recorded institutional process. The second asks about access to that process and must investigate potential claimants who are absent from the registry. The latter cannot be settled by greater precision within the registered population alone.

3.2 Observation, classification, and variable construction

The information produced from an admitted domain typically undergoes several operations: observation, recording, interpretation, variable extraction, and aggregation. A statement such as “three claims were resolved” may arise from a case-management database whose categories distinguish completed administrative entries, enforceable remedies, voluntary settlements, or procedural dismissals. The numerical result is meaningful relative to that taxonomy. An institutional inquiry should therefore recover how the recorded variable was produced, which distinctions were available to the person entering the data, and whether subsequent users retained the relevant qualifications.

To distinguish variable extraction from later reduction, let $V_\nu : D \rightarrow \mathcal{X}_\nu$ denote a declared feature map determined by a coding convention ν . The representation U introduced in Section 2 may be decomposed into feature construction and a further reduction G_ν when the model is designed in this way. The decomposition is recorded in Equation 4.

$$U_\nu = G_\nu \circ V_\nu, \quad V_\nu : D \longrightarrow \mathcal{X}_\nu, \quad G_\nu : \mathcal{X}_\nu \longrightarrow \mathcal{T}. \quad (4)$$

Equation 4 describes a stipulated construction, not a factorization available for every existing model. It permits distinct investigations of information discarded during variable construction and information combined by the subsequent reduction. If two admitted histories already receive the same feature vector despite differing on a question, later operations on that vector cannot recover the distinction without supplemental information. Conversely, a sufficiently informative feature vector may still be rendered inadequate by a later public summary.

The scope of a variable also depends on the stability of its measurement and interpretive conventions. A recorded category of “community participation” could refer to attendance at a meeting, an opportunity to submit comments, effective authority to revise a proposal, or participation in a binding vote. Different organizations may collect these indicators under the same name. The coding convention ν should specify its evidence sources, interpretive rules, intended population, recognized uncertainty, and permitted uses. Technical documentation proposals such as Gebru and colleagues’ *Datasheets for Datasets* call for descriptions of dataset motivation, composition, collection, and intended uses [18]. Their proposal supplies a precedent for documentation practices; the present paper’s specific obligations and variable-selection tests remain an exploratory extension.

Variable selection further depends on what a modeler considers an admissible question. A numerical indicator of compensation may answer an accounting question while leaving open who had authority to negotiate, which harms were eligible for compensation, and whether the compensation restored affected capacities. These are different institutional questions, some of which require interpretive rather than

scalar variables. Their practical relevance should be assessed before the coding convention is treated as sufficient for consequential use.

3.3 The boundary of observational identifiability

A formal distinction clarifies the difference between information loss within a domain and lack of identification caused by the admission process itself. Let $\mathcal{E}$ denote a declared set of candidate institutional histories for a particular research comparison. This set is an analytical supposition, rather than a census of an independently fixed historical universe. Let $R_\pi : \mathcal{E} \rightarrow \mathcal{O}$ represent everything that an institutional process makes available to the analyst, including its admission rule and subsequent record construction. Let $J_q^* : \mathcal{E} \rightarrow A_q$ represent an interpreted answer to a declared question q , assuming that the answer is defined under the stated comparative hypotheses.

The observation process may combine admission and coding, and it may omit an entire class of events or claimants. Its role differs from the reduction $U : D \rightarrow \mathcal{T}$ of Section 2, which is defined once a domain has been admitted. A necessary and sufficient condition for exact identification of J_q^* solely from observed output is the familiar fiber condition shown in Equation 5.

$$R_\pi(e) = R_\pi(e') \implies J_q^*(e) = J_q^*(e') \quad (e, e' \in \mathcal{E}). \quad (5)$$

Equation 5 is an instance of the elementary factorization criterion in Equation 2, applied to a broader observation process. Its value is conceptual: specifying the admission mechanism makes clear why improving an aggregate statistic over observed cases cannot necessarily answer a question concerning absent cases.

Proposition 3.1 (Admission-dependent non-identifiability). *Assume that $\mathcal{E}$ contains two candidate histories e_a, e_b with $R_\pi(e_a) = R_\pi(e_b)$ and $J_q^*(e_a) \neq J_q^*(e_b)$. Under these assumptions there is no function $F : R_\pi(\mathcal{E}) \rightarrow A_q$ whose composition with R_π equals J_q^* on $\mathcal{E}$.*

Proof. If such an F existed, equality of the two observed outputs would imply equality of their images under F . Those images would equal $J_q^*(e_a)$ and $J_q^*(e_b)$, contradicting the assumed difference. $\square$

Proposition 3.1 is an elementary observation about functions, not an original statistical identification theorem. Its use in this paper is to localize a particular representational limitation. It does not establish that either candidate history occurred, that the institutional process acted wrongfully, or that researchers have an unconditional right to obtain further records. A researcher must independently establish the plausibility of the candidate histories, the semantics of J_q^* , and the conditions under which additional evidence could be collected. Likewise, a model defined only on the registered population may perform well for a registered-population question while being insufficient for a broader question.

The conclusion remains informative in probabilistic settings only after specifying relevant observation models and uncertainty assumptions. Claims of statistical generalization require attention to selection mechanisms, comparability, estimation error, and potential dependence between admission and outcomes. The deterministic proposition supplies a limiting illustration, not a substitute for the statistical theory of missing data or for causal inference.

3.4 Constructed comparison of omitted claim histories

Consider a constructed watershed programme with two channels for environmental complaints. A digital portal receives reports from registered residents, while a separately administered local council hears claims from households lacking the documents required by the portal. A public performance dashboard draws exclusively on the portal and reports a stable percentage of “resolved” complaints. Suppose that the dashboard’s figures are identical under two candidate histories. In the first, the local council has effective means of securing institutional review. In the second, complaints received by the council are acknowledged informally and never enter a competent review procedure. The dashboard is stipulated to be identical in the two situations; no actual watershed, public authority, or affected community is identified by the example.

Table 3: Admission and variable-selection diagnostics in the constructed watershed comparison.

Analytical level	Possible loss or restriction	Additional evidence or authorized response
Population admission	Residents outside registration criteria disappear from dashboard denominators	Sampling-frame analysis, confidential intake review, and accessible routes to challenge eligibility
Observation and recording	Council complaints remain unrecorded or inaccessible	Independent record comparison under appropriate permissions and security safeguards
Variable coding	“Resolved” merges procedural closure and effective remediation	Category definitions, underlying case files, and review of downstream decision criteria
Question authorization	Model judged solely on portal throughput although substantive access is disputed	Justification of intended use, affected-party consultation, and explicit secondary questions
Institutional reliance	Aggregate dashboard substitutes for evidence about review rights	Competence audit, decision provenance, and comparative assessment of plausible alternatives

For a question about the fraction of portal records marked resolved, the dashboard may be sufficient, assuming reliable portal entries. For a question about effective opportunities for review among people exposed to environmental impacts, the two candidate histories yield different answers. Proposition 3.1 explains why the second question cannot be identified from the portal dashboard alone. The example concerns the outer admission boundary as well as the later summary map; restoring additional decimals to the dashboard would leave that boundary unchanged.

The contrast also illustrates how exclusion may occur at several levels. The separate council’s activity could be inaccessible to the analyst because its records are confidential. The council could use a classification incompatible with the portal schema. An institutional rule could deny the council competence to initiate a formal review. An investigator could simply define the study population as registered portal users. These possibilities cannot be collapsed into a single empirical explanation. Their significance depends on the research purpose, institutional authority, and the evidence available for adjudication.

Table 3 distinguishes the observational failures that the hypothetical case permits researchers to investigate. The table is an audit heuristic and does not claim that a particular remedy follows uniquely from any category.

The three columns of Table 3 should be read separately. Identifying an information loss is an analytical task. Establishing a reliable evidentiary correction is a methodological and practical task. Deciding who may disclose or compel further information is a legal and institutional task. Ethical evaluation requires an additional argument about affected interests, standing, confidentiality, and the relative burdens of corrective measures.

3.5 Question selection and interpretive authority

A declared question family $\mathcal{Q}$ is itself a selection device. A model can preserve information relevant to an agency’s efficiency target while excluding questions about who could contest a decision, which contributors were compensated, or what ecological conditions were damaged by an apparently successful intervention. A complete proof of adequacy relative to such a question family would remain silent about claims never submitted to the evaluation. Accordingly, the authority to choose $\mathcal{Q}$ needs an account distinct from the mathematical properties of the selected questions.

The distinction between procedural participation and substantive normative validity is especially important. A consultation may broaden the questions an investigator recognizes, but it may also reproduce established hierarchies of language, representation, or credibility. Fricker’s account of epistemic injustice provides relevant philosophical background for investigating the standing of persons as knowers

[16]. D'Ignazio and Klein's *Data Feminism* similarly foregrounds inequalities of power in the creation and use of data [9]. These sources motivate scrutiny of question formation. They do not determine which particular questions should have authority in a given legal or institutional setting.

One possible procedure is to maintain a *declared question register*. Each entry would identify the purpose for which a question was adopted, its affected population, the decision whose justification it informs, its interpretation requirements, and a means of recording objections to its exclusion or formulation. A change in the question register should not be presumed to describe a simple progression toward better knowledge: the changed questions may reflect new evidence, altered legal competences, disagreement among affected subjects, or shifting public priorities. The record should make these grounds traceable rather than presume one privileged historical direction.

Questions may themselves be incompatible. One group may request a public, geographically detailed disclosure of environmental damage; another may reasonably seek limits on disclosure because the same spatial resolution exposes households to retaliation. A declared question family can therefore contain competing demands that cannot be settled by maximal information retention. The ethical decision concerns authorized access, differentiated disclosure, and reviewable limits, alongside the preservation properties of the representation.

3.6 Protected evidence and constrained refinement

Increasing descriptive granularity may improve the ability to distinguish claims while increasing privacy, surveillance, or security risks. A proposed refinement should therefore specify which additional information is necessary for the intended task, which actors need access, and which forms of disclosure remain restricted. Nissenbaum's context-sensitive account of information flows provides a relevant point of comparison [33]. The paper's proposed distinction is between a representation's *capacity to preserve an answer* and an institution's *permission to reveal the evidence* on which that answer depends.

For the watershed example, a confidential reviewer might be authorized to determine whether a household was eligible to obtain review, while a public report publishes only aggregate outcomes accompanied by a meaningful description of their coverage limitations. The public report may remain deliberately coarse. The institutional pathway would nevertheless preserve a protected route to investigate exclusions or erroneous classifications. This is a candidate design, not an established universal standard. Its appropriateness depends on the reviewer's independence, the reliability of data custody, plausible retaliation risks, and the legal rights of persons seeking information or remedies.

The W3C PROV-O specification supplies a vocabulary for describing provenance relations among entities, activities, and agents [47]. An audit could employ that vocabulary to record how a dataset was produced, transformed, and authorized. A provenance trail does not compel access, validate a claimant's interpretation, or itself establish that an omitted party has standing. These tasks require different procedures and standards of justification. The same separation applies to dataset documentation: Gebru and colleagues' datasheet proposal supports the communication of collection choices and intended uses, while further institutional mechanisms may be necessary to respond to objections that documentation alone reveals but cannot correct [18].

A refinement can also introduce new exclusions by requiring forms of evidence unavailable to the relevant population. For example, adding a geolocation requirement to distinguish claims spatially could favor households with reliable addresses and mobile access. Corrective design should therefore assess changes to admission rules together with improvements in variable resolution. A useful improvement to statistical precision may coexist with a deterioration in procedural accessibility.

3.7 Rival explanations and empirical adjudication

The analytical distinctions proposed here can guide research, but a study must distinguish a representational effect from competing institutional explanations. Suppose the watershed programme reports fewer successful appeals from communities omitted from the portal. That pattern could arise because excluded

claims were absent from the dashboard, because a separate legal rule limited competence, because remediation resources were scarce, or because recorded claims differed in relevant ways. The model’s construction may contribute to one or more mechanisms, but a difference in observed outcomes alone does not isolate its effect.

An empirical design should reconstruct the sequence or partial ordering of admissible decisions, record-entry practices, administrative classifications, and implementation evidence, while keeping evidential timestamps distinct from event precedence. Comparisons might involve the same institution under alternative admission rules, parallel institutions with similar mandates and different intake channels, or documentary analysis of a coding revision. Each comparison requires its own assumptions about case comparability, endogenous policy change, and relevant confounders. The research cannot treat an apparently successful refinement as proof that information loss had caused the prior outcome.

The framework permits meaningful negative cases. If independent review establishes that the portal captured all relevant claimants within a justified jurisdiction, a broader exclusion hypothesis would be weakened for that setting. If a proposed variable refinement leaves relevant legal competences unchanged and does not affect decisions, the explanatory contribution of that refinement would be limited. If a conventional qualitative case study recovers the omitted claims with equal reliability and lower cost, the formal diagnostic may add little practical value. These observations should be reported as constraints on the framework, not reclassified automatically as further evidence of systemic loss.

An audit programme may consequently proceed through a bounded sequence: define the claimed use of a representation; identify its admission, variable, and question-selection rules; record who authorized them; seek independently grounded counterexamples and alternative explanations; evaluate privacy and evidentiary constraints; and compare feasible procedures for correction. The order describes a practical research workflow and carries no claim about a universal temporal ontology. In cases with incomplete or disputed records, the appropriate outcome may remain an explicit finding of indeterminacy.

3.8 Scope of the proposed diagnostic

The main result of Section 3 is a structured distinction among domain admission, variable construction, question selection, and consequential reliance. The formal condition in Equation 5 and Proposition 3.1 show why questions about omitted histories can be unidentifiable even when a model accurately summarizes its admitted observations. They establish a restricted property of a stipulated function model. Neither determines how to identify a legitimate claimant, what evidence may be collected, or how responsibilities should be allocated when institutional interests conflict.

The approach also relies on historically situated choices about the candidate comparison space $\mathcal{E}$. If evidence or interpretive concepts make additional histories intelligible, the model can be revised; a previously observed agreement does not establish universal sufficiency. Normative disagreement can persist over the question itself, not only over the data needed to answer it. Informational completeness, procedural legitimacy, and moral justification remain separate assessments. Section 4 examines preservation of declared normative questions under reduction and the information losses arising under specified institutional constraints.

4 Normative Information and Preservation Conditions

The analytical task of this section is to distinguish the information required to answer a normative question from the information that an institution is warranted in collecting, retaining, or disclosing. Section 2 introduced query-relative factorization; Section 3 examined how the admitted population and coding conventions delimit the information available to that factorization. The present section develops a policy-sensitive account of preservation, approximation, and disclosure. It first treats changes in the questions to which an effective description is held accountable, then derives an elementary lower bound for the error of a quantitative summary, and finally examines the possibility that protected access can preserve a route to review without enabling unrestricted public disclosure. The argument employs finite or otherwise

explicitly bounded comparison spaces. Its formal conclusions identify compatibility and information limits; the legitimacy of the questions and of access permissions requires further normative reasoning.

4.1 Preservation relative to a declared family of questions

Let D be the admitted history domain, $U : D \rightarrow \mathcal{T}$ the representation used for a specified institutional purpose, and $\mathcal{Q}$ a declared family of questions. For every $q \in \mathcal{Q}$, the interpreted answer function $J_q : D \rightarrow A_q$ is taken to be total and deterministic within the stipulated comparison. These assumptions isolate one issue of representation adequacy. An actual institutional question may instead admit multiple defensible answers, incomplete evidence, or a legally imposed rule that excludes the question from a particular body's jurisdiction. Such cases can be modeled only after the distinct grounds for uncertainty or exclusion are specified.

For each map $F : D \rightarrow X$ define its observational equivalence relation $\sim_F$ by $h \sim_F h'$ when $F(h) = F(h')$. The query family similarly induces a relation $\sim_{\mathcal{Q}}$ when every declared answer agrees. The exact preservation condition is expressed in Equation 6.

$$\sim_U \subseteq \bigcap_{q \in \mathcal{Q}} \sim_{J_q} \quad \text{on } D. \quad (6)$$

A representation satisfying Equation 6 permits each declared answer to factor through the attained descriptions $U(D)$. The statement follows directly from the fiberwise criterion in Equation 2. It expresses a property of functions on the specified domain, without ranking the relative moral importance of the questions. An administrative description can preserve every question specified by an agency while failing to record claims that affected communities regard as essential. Conversely, a description may preserve more distinctions than a particular administrative report requires, and those additional distinctions may appropriately remain unavailable to some audiences.

The choice of $\mathcal{Q}$ is therefore part of the authority problem. List and Valentini examine the possibility of coarser-grained political normative facts, which provides an important philosophical precedent for asking which descriptive levels support normative inquiry [29]. The present analysis focuses more narrowly on the duties of those who authorize the questions and construct the corresponding descriptions. A preservation proof offers no independent justification for the answer functions that it preserves. The question register proposed in Section 3.5 should accordingly record competing formulations, the interpretive grounds for each, the populations whose circumstances they address, and the procedures by which their selection can be challenged.

An effective description may also be evaluated under a family of interpretations $\mathcal{I}$. For each admissible interpretation $i \in \mathcal{I}$, a question may have an interpreted answer $J_{q,i}$. A robust preservation requirement replaces $\mathcal{Q}$ in Equation 6 by the explicitly enumerated collection of question–interpretation pairs. This is stronger than agreement under one privileged interpretation. The requirement does not resolve which interpretations are ethically defensible: admitting an interpretation to methodological comparison can be justified by evidence, plausible standing, or unresolved disagreement, while its legal authority and moral adequacy remain separate matters. Treating incompatible interpretations as distinct indexed questions avoids the appearance that a single summary has adjudicated their substantive conflict.

4.2 Revision of questions and represented populations

The adequacy of a model is sensitive to the domain and question family at which it is evaluated. Suppose an institution first assesses U against $\mathcal{Q}_0$ and later evaluates it against $\mathcal{Q}_1$ with $\mathcal{Q}_0 \subseteq \mathcal{Q}_1$. If U preserves $\mathcal{Q}_1$, it also preserves $\mathcal{Q}_0$. The converse has no general warrant. The logical inclusion is immediate from Equation 6; the substantive significance depends on what motivated the additional questions. A wider inquiry might recognize previously disregarded forms of harm, respond to a different institutional mandate, or introduce newly contested standards of evidence. The fact of enlargement conveys no automatic claim of ethical progress.

A different problem arises when the admitted population expands from D_0 to D_1 , where $D_0 \subseteq D_1 \subseteq \mathcal{H}$. A representation that was sufficient on D_0 may identify two newly comparable histories on D_1 while a declared question distinguishes them. Even if a coding scheme retains the same variable names, the extension can change which claims its summary is capable of answering. The resulting failure should be distinguished from changes in the interpretive function $J_{q,i}$ itself. Domain enlargement, query enlargement, revision of the answer semantics, and changes in allowed disclosure are four analytically different modifications to a representational policy.

This distinction becomes consequential in historically contested institutional settings. A historical claim can receive a different recognized meaning when the relevant community, competent authority, or background account of an obligation is revised. The paper treats such revisions as changes in specified relational and evidential structures; it presupposes neither a total event order nor a universal division between past and future. The effect of a change must be established for the relevant histories and questions. The methodological analogue in responsible modeling is the need to revisit assumptions and uncertainties as a model is used outside its initial setting [42]. That literature motivates checking the conditions of reuse; it does not imply that every amendment is beneficial or that any fixed model has an unrestricted domain of validity.

A practical maintenance rule is to document a preservation claim together with its domain, coding convention, question register, interpretation versions, and validation evidence. When any one of these changes, the claim can be re-tested. The re-test may conclude that the available evidence is too sparse to resolve the effect. It may also conclude that the original representation remains adequate. Both outcomes are legitimate results of a bounded inquiry and should be recoverable from the provenance record.

4.3 Error bounds for quantitative answers

Exact preservation may exceed the needs of some quantitative decisions. An administrative model may need to estimate a count or a measured exposure within an explicit tolerance. A tolerance is meaningful only when the answer space carries a justified error metric, and when the consequences of different errors have been assessed independently. Procedural standing, a legal right to review, or the existence of a protected claimant should not be reduced by default to a numeric approximation problem.

Consider a finite nonempty domain D , a real-valued question $J : D \rightarrow \mathbb{R}$, and an effective description $U : D \rightarrow \mathcal{T}$. For each attained coarse state $t \in U(D)$, the U -fiber has a minimum and maximum value of J . The worst possible within-fiber range is defined in Equation 7.

$$\text{osc}_U(J) = \max_{t \in U(D)} \left(\max_{h \in U^{-1}(t)} J(h) - \min_{h \in U^{-1}(t)} J(h) \right). \quad (7)$$

The range in Equation 7 depends on the chosen history domain and on the interpretation of the measured quantity. It represents descriptive ambiguity, not a moral cost. The following elementary proposition identifies the best uniform accuracy possible when an estimator has access solely to $U(h)$.

Proposition 4.1 (Uniform approximation from a finite reduced record). *For nonempty finite D and real-valued J , the minimum worst-case absolute error among functions $a : U(D) \rightarrow \mathbb{R}$ is one-half of the maximum within-fiber range. In particular, the equality in Equation 8 holds.*

$$\min_{a: U(D) \rightarrow \mathbb{R}} \max_{h \in D} |J(h) - a(U(h))| = \frac{1}{2} \text{osc}_U(J). \quad (8)$$

Proof. For any attained t , take histories on that fiber with the minimum and maximum values of J . By the triangle inequality, any single number $a(t)$ has absolute error at least half their difference on at least one such history. Taking the maximum over the fibers gives the lower bound in Equation 8. Assigning $a(t)$ the midpoint of the minimum and maximum on each fiber achieves precisely half that fiber's range. The maximum error therefore equals the stated bound. $\square$

Proposition 4.1 is a basic minimax property of finite partitions. It entails no new decision theory and makes no claim that a midpoint estimate is institutionally permissible. When an error of a certain magnitude can impose qualitatively different burdens on different populations, a uniform absolute-error criterion may be inappropriate. Stochastic observations, interval-valued evidence, asymmetric error costs, and contested measurement conventions likewise require other models. Established work comparing the informativeness of statistical experiments offers a distinct mathematical tradition for examining information in relation to decision problems [5]. The finite partition result here is considerably narrower and should not be confused with Blackwell’s comparison of statistical experiments.

A coarse report might, for example, estimate mean exposure adequately for a planning exercise while obscuring which household has an actionable claim against a particular institution. Low aggregate error would then coexist with a loss of eligibility information. The responsibility of the model’s users depends on the decisions attached to each output. A numerical error threshold alone cannot determine whether the representation gives affected people an effective route to review.

4.4 Preservation and restricted disclosure

A model may be constructed for multiple audiences whose legitimate access conditions differ. An internal reviewer may require identifiable testimony to assess procedural eligibility, whereas a public dashboard may appropriately disclose only region-level aggregate quantities. Under such an arrangement, the institution can possess information that the public interface is not authorized to reveal. The choice is a governance matter involving privacy, evidential rights, accountability, and potential coercion, rather than an inherent virtue of descriptive fineness.

Let $U : D \rightarrow \mathcal{T}$ denote a protected representation and let $S_a : \mathcal{T} \rightarrow \mathcal{Z}_a$ be the information released to an audience a . For each audience, define a policy-imposed equivalence relation $R_a \subseteq D \times D$: any two histories related by R_a must be indistinguishable through that audience’s permitted channel. A channel complies with this simplified constraint when $hR_a h'$ implies $S_a(U(h)) = S_a(U(h'))$. This relation is a stipulated confidentiality rule whose ethical and legal appropriateness must be established separately. Some real disclosure systems use randomized or risk-based protections, which cannot be captured by this deterministic condition.

A direct tension arises when an audience’s requested answer separates histories that the disclosure policy requires to remain indistinguishable.

Proposition 4.2 (Disclosure conflict for a declared question). *Let q be a question with a deterministic answer function $J_q : D \rightarrow A_q$. Suppose there exist $h, h' \in D$ such that $hR_a h'$ and $J_q(h) \neq J_q(h')$. No deterministic disclosure channel satisfying the stipulated R_a -indistinguishability rule can allow J_q to be reconstructed exactly from its disclosed output for every history in D .*

Proof. The policy requires the channel to give the same public output at h and h' . An exact reconstruction function would then give the same answer for both histories, contradicting $J_q(h) \neq J_q(h')$. $\square$

Proposition 4.2 establishes a compatibility limit relative to a chosen confidentiality relation. It does not establish that a public question should be prohibited, that a particular secrecy rule is legitimate, or that confidentiality inevitably defeats accountability. Access can be differentiated by role, procedure, or protected review. Nissenbaum’s account of context-dependent information-flow norms provides the ethical context for assessing those choices [33]. Provenance standards can also record who created or transformed an output, though such records alone do not confer a right of access or an ability to compel correction [47].

The implication for institutional design is a separation of three functions: whether the question can be answered from the protected record; whether an authorized reviewer can evaluate it; and whether the answer or supporting evidence may be disclosed to the broader audience. An apparently opaque public summary can coexist with meaningful protected contestation if the review mechanism has sufficient competence, practical accessibility, independence, and authority to secure consequences for its findings. Conversely, a public aggregate accompanied only by the promise that sensitive details exist elsewhere

Table 4: Distinct judgments in policy-sensitive preservation analysis.

Judgment	Question assessed	Evidence or warrant required
Exact preservation	Is the answer constant on every attained reduction fiber?	Declared domain, coding map, answer semantics, and loss witnesses
Quantitative tolerance	How much worst-case answer variation does reduction permit?	Justified metric, bounds, consequences of error, and model coverage
Permitted disclosure	May the relevant information circulate to the named audience?	Contextual privacy duties, access permissions, and independent review routes
Avoidability of loss	Is a feasible alternative able to preserve the relevant distinction?	Comparison of permissible alternatives, costs, institutional capacities, and risks
Ethical adequacy	Are the chosen questions, exclusions, and uses defensible?	Affected-subject accounts, normative reasoning, legal competence, and evidence of consequences

may offer no meaningful safeguard. These are hypotheses about institutional performance requiring observation, not consequences of the two propositions alone.

4.5 Avoidable loss under institutional constraints

A claim that information was unavoidably lost must name the alternatives among which the institutional choice was made. Let $\mathfrak{F}_a$ denote a set of feasible and permitted descriptions or disclosure channels for a declared decision setting and audience a . Its membership may depend on lawful competence, confidentiality conditions, measurement costs, security requirements, and the rights of affected persons. A loss of the answer to q is *avoidable relative to $\mathfrak{F}_a$* when some admissible alternative preserves J_q on the declared domain, and *unavoidable relative to $\mathfrak{F}_a$* when none does. Neither label has an unqualified meaning outside the stated feasible set.

This definition prevents a common ambiguity. A model designer may call a distinction technically irrelevant because the presently selected variables cannot retain it. A different feasible design may preserve that distinction at moderate cost. Conversely, a reviewer may demand a finely identifiable public record whose disclosure is prohibited on independently justified grounds. Calling its absence a representational failure without acknowledging those grounds would conflate information quality and permissible circulation. Establishing the feasible set requires empirical knowledge of infrastructure and resources, examination of binding law where applicable, and an explicit normative account of affected rights. A mathematical comparison can organize those claims, but cannot certify them from the mapping functions alone.

Table 4 records a compact set of distinctions for evaluating alternative arrangements. Each row identifies a different question and a correspondingly different class of evidence. The table is a proposal for transparent analysis, rather than an established audit standard.

Table 4 also points to cases of permissible indeterminacy. An institution may reasonably decline to certify a claim when available records do not support it, while explaining the limits of that determination and maintaining lawful channels for additional evidence. The quality of such a decision cannot be inferred from a low numerical information-loss score alone. Transparency about uncertainty can be more responsible than presenting a precise output whose interpretation exceeds the supporting records.

4.6 Constructed comparison of protected review procedures

Consider a hypothetical coastal monitoring programme whose public dashboard reports the total number of claims reviewed within each coastal district. Its protected case file records whether each individual claim concerns physical exposure, ecological damage, or a decision about procedural standing. For the present comparison, suppose two records yield identical district totals: in history h_1 , an affected household was eligible to seek review; in history h_2 , an administratively equivalent complaint was refused

because the household's residence was not recognized by the intake rule. The public report cannot distinguish h_1 from h_2 . A question about district workload is preserved, while the eligibility question is lost from that interface.

Three policy designs are relevant. Under design A, the institution publishes aggregate counts and offers no independent procedure for investigating exclusion. Under design B, a legally authorized reviewer can inspect protected case records and issue a reasoned determination whose implementation is monitored. Under design C, the public release exposes precise household location and claim status, enabling independent identification at the cost of potentially serious confidentiality risks. No ordering of these designs follows from coarse-graining theory alone. Design B may be preferable for specific rights and risks if the reviewer is accessible, genuinely independent, and able to affect institutional decisions. Design C may be unlawful or disproportionate. Design A may be adequate for a narrowly bounded statistical publication while being inadequate as the sole evidentiary basis for a claim that all affected households received access to review.

The example also reveals a possible inferential error. A later increase in successful appeals under design B would not, by itself, establish that protected information access caused the change. Legal competence, resource availability, staffing, changes in intake categories, and shifts in the affected population could also matter. Empirical evaluation would reconstruct decision records and comparison groups where feasible, document selection into review, and examine unsuccessful as well as successful claims. Such investigation requires evidence beyond the stipulated histories; the formal model supplies questions to ask, not the causal answer.

The design choices also depend on whose interpretations enter the question register. A monitoring body could prioritize immediate physical exposure, while coastal residents emphasize the historical continuity of livelihood and ecological relations. Incorporating the latter may alter admissible histories, coding conventions, and the competence required for effective review. As classification scholarship emphasizes, institutional categories acquire their practical meaning through systems of work and governance [7]. The example illustrates the need to examine those systems rather than assuming that a mathematically finer variable automatically constitutes a more just representation.

4.7 Evidential and ethical scope of preservation claims

The analysis in Section 4 provides bounded preservation diagnostics for a declared domain and question family, an elementary worst-case error result for finite quantitative settings, and a formal incompatibility test for stipulated confidentiality requirements. These contributions concern the information available through a description or a channel. They do not fix which historical accounts should be admitted, which audiences possess legitimate authority, or which losses count as acceptable under the rights and obligations at stake.

The principal methodological implication is the need for a *separable warrant record*. Formal claims require definitions and proofs or counterexamples. Empirical claims require observations relevant to the represented population and independently assessed rival explanations. Legal claims require competent jurisdiction-specific interpretation. Ethical claims require arguments sensitive to the standing, agency, dependency, and vulnerability of affected persons and communities. A concise representation is responsible when its purpose and limits are intelligible, its governance can be contested where appropriate, and consequential reliance remains proportionate to its demonstrated adequacy.

Two limits remain particularly important. First, for an open-ended family of possible interpretations or future queries, the present finite diagnostics cannot demonstrate universal preservation. Second, enriching a record can change the historical and institutional relations under examination by changing practices of observation, classification, or surveillance. Its effects therefore require inquiry into the actual consequences of collecting and circulating the information, including unintended burdens. Subsequent sections address the ethical significance of historically discarded annotations and the mechanisms through which representational exclusions become epistemic injustices. The more general question of how communities and institutions can challenge representational authority will remain central throughout

the paper.

5 Historical Structure and Discarded Annotations

This section examines the ethical and epistemic consequences of discarding information about event precedence, the changing recognition of claims, institutional continuity, and potentially irreversible harms. Its objective is to specify which historical distinctions can be lost even when a reduced representation preserves the aggregate quantities for which it was designed. The discussion first defines several orders and annotations that must be kept analytically distinct, develops an elementary order-loss criterion with a counterexample, and then examines institutional succession, ecological thresholds, and provenance-sensitive review. The resulting framework provides a method for identifying losses relevant to a declared question. The selection of the question, the admissibility of historical evidence, and the legitimacy of resulting institutional obligations require further arguments.

5.1 Event precedence and the limits of chronological reconstruction

An institutional history may be represented provisionally by a finite enriched record $h = (E_h, \prec_h, \alpha_h, \kappa_h, \omega_h)$, where E_h is a set of identified events, $\prec_h$ is a strict partial order of event precedence, α_h is a family of recorded annotations, κ_h represents documented institutional competence and review arrangements, and ω_h contains stipulated outcome observations. The tuple is a modeling device, rather than an inventory of the constitutive elements of reality. Its components refer to events and relations that have been identified under declared evidential and interpretive conventions. Persons and institutions are reconstructed from relational patterns, records, practices, and recognition, and are not stipulated as elementary event vertices. Where a two-complex representation is used, vertices may index events and edges couplings, while the interpretation of faces remains specific to the application.

A strict partial order is irreflexive and transitive. It admits incomparable event pairs and requires no privileged present, universal clock, or total temporal ordering. The use of a partial order recalls Lamport's technical analysis of event order in distributed computing [27]. That technical result supplies a precise mathematical precedent, while the choice of socially relevant events and the evidential warrant for social precedence remain problems of historical interpretation. The order $\prec_h$ can express stipulated or documented precedence relations without treating them automatically as causal dependencies. Additional evidence is required to infer that one action changed another action's occurrence, meaning, or consequences.

Three further distinctions are especially important. A *recording order* describes when an archive received or created evidence. An *evidential order* describes justified dependencies among sources or interpretations. A *narrative order* describes the sequence through which an account presents events. A document recorded later may provide evidence concerning an event represented as earlier; a narrative may reorder its presentation to clarify an argument; and two apparently successive events may have no established causal relation. An audit should therefore specify the kind of ordering that underwrites each historical claim. A timestamp without an account of its source, resolution, and relation to the claim provides limited grounds for adjudicating precedence.

Institutional path dependence introduces another form of historical sensitivity. Pierson's analysis of increasing returns emphasizes that timing and sequence can be consequential in political processes, including circumstances where established arrangements become costly to reverse [39]. The present section draws on that insight as a methodological motivation. A historical dependence claim requires a mechanism linking an earlier configuration to subsequent institutional possibilities, including evidence capable of distinguishing that mechanism from stable background conditions, correlated events, and retrospective storytelling. Sequence sensitivity alone establishes neither a unique causal path nor an obligation of repair.

5.2 Recognition, standing, and consequential annotations

The annotation family α_h may contain several kinds of structured information: the stated basis for admitting a complaint, an evidential qualification attached to a recorded contribution, the scope of an authority's competence, a protected assertion of ecological harm, or an institutional decision concerning a claimant's standing. These annotations are historically situated and may themselves be contested. An annotation is accordingly modeled as a claim with a provenance and warrant, rather than an automatically true property of the underlying events.

For example, a complaint record can indicate that an affected subject was recognized by a particular office, that the office had authority to investigate, and that an appeal route was available. Another record may preserve the occurrence of a complaint while discarding these qualifications. An aggregate count of complaints can be correct in both cases. The two representations nevertheless support different questions about effective participation, admissibility, and institutional response. The argument extends the classification concern developed by Bowker and Star [7]; it does not attribute the present formal distinction to their work.

The identity and continuity of an affected subject also require evidential interpretation. A community, cooperative, or regulatory body may change membership, legal form, or jurisdiction while a claim concerning its historically constituted relations remains relevant. Institutional records may support a partial correspondence between recognized claimants in different configurations. Such correspondences need not be functional, unique, or universally accepted. The retention of a stable administrative identifier cannot alone establish moral continuity, while a changed identifier does not suffice to establish that an obligation has ended.

The representational policy should record the particular institution competent to assess these issues, its jurisdictional basis, the status of affected participants, and any procedures for correcting its classifications. This is related to Iris Marion Young's distinction between liability-oriented and structural accounts of responsibility [50]. The distinction supports attention to responsibility across continuing social processes, while the allocation of any concrete duty in a particular case depends on independent legal and normative argument.

5.3 A criterion for preserving order-sensitive questions

Let D be a nonempty domain of enriched histories, $U : D \rightarrow \mathcal{T}$ an effective description, and $J_{\text{ord}} : D \rightarrow A$ a declared order-sensitive question. The query may concern whether a specified review decision had recognized competence at the point relevant to a subsequent intervention, provided that both the competence criterion and the event relation have been stipulated. Equation 9 states the condition for recovering the answer from the reduced record.

$$\forall h, h' \in D, \quad U(h) = U(h') \implies J_{\text{ord}}(h) = J_{\text{ord}}(h'). \quad (9)$$

Equation 9 is the ordinary fiberwise factorization criterion applied to an explicitly historical question. The following proposition identifies a direct form of failure under event-order erasure.

Proposition 5.1 (Failure of order-sensitive recovery). *Suppose D contains histories h_1 and h_2 with identical retained event labels, annotation totals, and outcome totals. Suppose that a reduction U depends only on those retained quantities, so that $U(h_1) = U(h_2)$. If a declared question satisfies $J_{\text{ord}}(h_1) \neq J_{\text{ord}}(h_2)$ because of an order distinction erased by U , there exists no function $\bar{J}_{\text{ord}} : U(D) \rightarrow A$ that recovers the declared answers on all of D .*

Proof. If such a function existed, both histories would receive the same value $\bar{J}_{\text{ord}}(U(h_1)) = \bar{J}_{\text{ord}}(U(h_2))$. Their stipulated question values differ, contradicting the required equality. The proposition is a direct consequence of Equation 9 and adds no new set-theoretic theorem. $\square$

Consider a constructed public-restoration programme with two identified events: a , a committee decision recording valid authorization under a stipulated rule, and b , a material intervention affecting a

local habitat. Both histories contain one recorded authorization and one intervention and report an identical aggregate restoration expenditure. In history h_1 , the authorized decision precedes the intervention. In history h_2 , the intervention precedes the decision. For the sole purpose of this example, suppose the applicable procedural rule requires authorized review to precede the intervention, and suppose the evidence adequately establishes the two orders. A report preserving only event counts and total expenditure identifies these histories even though the declared procedural question distinguishes them.

The example is limited in three ways. First, actual procedural validity depends on the governing legal framework and relevant exceptions; the constructed rule serves only as an explicit premise. Second, precedence does not independently establish informed consent, substantive justice, or ecological consequences. Third, the distinction may already be recorded in a separate audit log available to competent reviewers. In the last case, the public aggregate can remain adequate for its stated descriptive purpose while the combined institutional information system supports the additional question. The relevant object of criticism is the entire policy of evidential access and consequential use, rather than the isolated compression operation.

5.4 Historical paths and the continuity of claims

The order-erasure proposition concerns a single defined question. Questions about claim continuity can require a richer historical reconstruction. A claimant may initially submit a complaint, receive an institutional classification, face a transfer to a new authority, and obtain a decision through a revised procedure. The order, competence, and interpretive status of these events may affect what the later institutional decision means. A compressed summary reporting a complaint and a disposition can leave the route to that disposition unrepresented.

A useful analytical object is a *claim path*, understood as a supported sequence or partially ordered configuration of claim-related events with links to relevant evidence and institutional jurisdictions. Distinct paths may share an administrative endpoint. Preserving the endpoint therefore provides weaker information than preserving the conditions of admissibility, transfer, and review. The distinction is especially important where recognition is contested or where multiple institutional authorities interpret the same claim in different ways. A later recorded disposition may correct, confirm, or narrow an earlier decision, and the procedural significance of each relation should be specified independently.

For a declared family of questions $\mathcal{Q}_{\text{hist}}$, the historical adequacy of a description depends on preservation of each J_q over the admitted domain. When the family includes the availability of a protected challenge, documentary continuity, and the reasons given for a refusal, a representation containing only the final classification may fail. The appropriate response could be to add a structured annotation, preserve an authorized linkage to protected evidence, or revise the institutional rule determining who can request review. The framework offers no presumption that every path should be disclosed publicly: confidentiality, safety, and legitimate restrictions on access can shape a differentiated preservation design.

Young's analysis of responsibility for historical injustice supports attention to continuing structural processes across changed institutional circumstances [50]. Historical significance, however, does not transfer automatically through a data schema. Whether a present institution inherits a specific duty can depend on institutional succession, existing law, the character of the harm, the capacities of affected participants, and contested ethical premises. A historical annotation may provide evidence relevant to these questions without deciding them.

5.5 Ecological thresholds and irreversible burdens

Historical information can matter for a further reason: a system's capacity for continuation may depend on changes that cannot be reversed within the relevant set of physically, technologically, and institutionally feasible interventions. The term *irreversible* consequently requires a stated system boundary, temporal or intervention horizon, uncertainty convention, and account of relevant capacities. It should not be inferred merely from a monotonically increasing administrative total.

The Intergovernmental Panel on Climate Change reports irreversible observed effects in certain ecosystems, including some extinctions attributable to climate change, and distinguishes these from other impacts that approach irreversibility under specified conditions [22]. That assessment provides empirical motivation for treating some historical ecological distinctions as consequential. It offers no case-specific evidence for the hypothetical programme considered here and establishes no universal threshold applicable to arbitrary institutional models.

Consider an institution publishing average habitat quality across several monitoring units. Two histories can have the same average quality at the assessment points while differing in whether one unit crossed a stipulated ecological-loss threshold between observations. The histories may also differ in the species, community dependencies, or restoration options affected. A mean value alone may therefore fail to answer whether a protected threshold was crossed or whether an available remedy can restore a particular functional capacity. Scientific assessment would require measurements, a validated threshold mechanism, uncertainty estimates, and a defensible relation between the measured variables and the ecological claim. A normatively consequential threshold also requires justification of its institutional relevance and treatment of affected subjects.

An information-preserving response can take the form of a protected threshold-crossing record, interval-specific exposure information, or a linked scientific assessment retained under competent review. Each design has potential costs: false positives, disclosure risks, exclusion of difficult-to-measure ecological relations, and a risk of conflating regulatory indicators with ecological reality. When an aggregate adequately preserves the ecological distinctions material to the declared inquiry, additional historical detail may provide little explanatory gain. The framework should report this outcome explicitly.

5.6 Provenance records and differentiated access

Historical annotation depends on provenance: the documented origin, processing, interpretation, and revision of evidence. W3C PROV-DM provides a domain-general vocabulary concerning entities, activities, agents, generation, use, derivation, and revision [46]. Such a vocabulary is useful for structuring an audit trail, although interoperability and provenance completeness do not guarantee truthful sources, lawful decisions, or ethically warranted interpretations. A provenance record may faithfully describe an unjust exclusion; it may also reproduce a contested classification without recording why the classification was accepted.

Table 5 proposes a compact research inventory connecting potential losses to specific questions, evidence types, and qualifications. The inventory distinguishes the adequacy of a public report from the adequacy of the protected institutional evidentiary pathway. Where identifiable claimants or sensitive environmental information are concerned, access controls may be essential to the legitimacy of the record itself [33]. Review permissions must therefore be evaluated jointly with retained information.

A minimum viable audit would first identify the consequential decision and its declared questions. It would then reconstruct the admitted history set, specify the retained and discarded annotations, trace available evidence with access permissions, compare plausible histories with identical published summaries, and determine whether competent review could distinguish those histories. The procedure permits a finding of genuine adequacy, documented loss, competing interpretation, or indeterminate evidence. None of these findings alone establishes liability or a uniquely preferred remedy.

5.7 Competing explanations and negative cases

Several rival mechanisms can explain why a published summary fails to answer a historical claim. A reporting system may lack the records because the underlying activity was never observed. A competent institution may restrict disclosure under a legitimate confidentiality rule. The relevant legal duty may turn on a different fact than the researcher expected. A scientific indicator may be insufficiently reliable to support the proposed threshold claim. An apparent path-dependent result may arise from persistent

Table 5: Historical information losses and corresponding evidential checks in a proposed representational audit.

Lost distinction	Question requiring the distinction	Evidence and qualification
Event precedence	Whether a stipulated review occurred within its competent procedural position	Linked event records; establish order and jurisdiction independently
Recognition status	Whether a claimant had standing and a route to challenge classification	Decision history and review rule; protect sensitive identities
Claim succession	Whether an obligation or remedy survived an institutional transfer	Succession instruments and transfer records; avoid assuming moral continuity
Ecological threshold	Whether a protected loss or exposure crossed a relevant boundary	Validated measurements and domain-specific science; retain uncertainty
Evidential revision	Whether a final summary incorporates challenged evidence and corrected interpretations	Provenance chains and reasons for correction; inspect unresolved disagreements

background factors rather than the order of events. Distinguishing these situations prevents the language of representational exclusion from becoming a generic label for every unresolved case.

The proposed diagnostic makes limited empirical predictions. For a specified question and domain, records that collapse histories with different justified answers will require additional evidence or stronger assumptions to support that question. Where an independently maintained evidentiary pathway restores the relevant distinction, the reduced public report need not impair institutional adjudication. Where the allegedly lost distinction has no effect on any authorized or defensible question, the case for a corrective refinement weakens. If archival review reveals that the supposed sequence difference was a recording artifact, the particular historical-loss claim fails. A research design should document these negative cases alongside positive findings.

Comparisons with conventional historiography, qualitative process tracing, legal document review, and ecological assessment are necessary. The use of a relational data model may improve clarity about which records support a claim, but it may also impose coding costs and interpretive simplifications. The measure of contribution is whether the additional formal interface improves accountable inquiry under specified institutional constraints, including access, confidentiality, error, and feasibility. No advantage over established methods is assumed.

5.8 Interpretive scope and transition to epistemic exclusion

Section 5 identifies a family of historical distinctions whose loss can be material for defined normative questions. Event order, recognition annotations, claim paths, and ecological thresholds have different evidential requirements, and each may depend on interpretations that remain disputed. The formal result in Proposition 5.1 states a bounded information limit. It does not justify the institutional rule, prove a causal relation, or settle a disagreement concerning the moral significance of a history.

The ethics of selecting and discarding annotations thus concerns both the preservation of available evidence and the authority to specify which claims can become questions for the representation. Section 6 will examine how testimonial credibility, shared interpretive resources, and classification conventions can determine whose accounts enter that question register at all. Its focus will be the institutional mechanisms of representational exclusion and the conditions of responsible correction.

6 Epistemic Injustice and Representational Exclusion

This section investigates the institutional processes through which a person's account becomes eligible for reception, interpretation, evidential assessment, and inclusion in an effective representation. Its objective is to connect the ethics of classification and epistemic injustice with identifiable decisions in

representational practice. The discussion distinguishes testimonial credibility, interpretive resources, formal admissibility, and subsequent aggregation; introduces a provisional model of the representational pathway; establishes a narrow limit on audits confined to its final output; and develops a constructed comparison with evidential alternatives and counterexamples. The inquiry concerns procedures through which testimony and claims acquire institutional standing. Diagnosing a representational failure still requires evidence about the relevant practices, the validity of the claim under examination, and the legal and ethical considerations governing institutional recognition.

6.1 Epistemic standing within representational practice

An institutional dataset records decisions about what can enter a description. Those decisions may occur before a statistical reduction is applied: a person may lack access to a complaint channel, an experience may lack a category recognized by the reporting system, or a statement may receive insufficient credibility for inclusion. Consequently, an accurately computed summary may preserve its admitted records while failing to represent claims that were filtered out at intake or interpretation. This is a question about the formation of the represented domain and the practices that mediate knowledge claims, beyond the narrower question of whether a mapping faithfully compresses its input.

Fricker’s analysis of epistemic injustice distinguishes wrongs concerning the credibility afforded to speakers and the interpretive resources available for understanding social experience [16]. The present paper uses that distinction to examine institutional records. Testimonial injustice is relevant when a speaker receives a credibility deficit connected to identity-based prejudice. Hermeneutical injustice is relevant when unequal participation in the creation of shared interpretive resources disadvantages a subject in making an experience intelligible. These philosophical concepts have demanding conditions of application. An absent data field, an inconclusive testimony, or an adverse judgment does not by itself demonstrate either form of injustice.

Dotson’s account of epistemic violence concerns testimonial encounters in which hearers, through pernicious ignorance, fail to meet speakers’ vulnerabilities; she distinguishes practices of silencing within that setting [11]. Her account permits more careful examination of interactions in which a formal invitation to testify coexists with conditions that prevent testimony from being meaningfully received. Medina’s analysis of epistemic resistance provides a further resource for considering structured ignorance, imposed silences, and the potential value of engagement with differently situated interpreters [30]. These works address matters of epistemic responsibility considerably richer than the formal loss functions employed here. Their contributions inform the identification of possible institutional mechanisms; they do not establish the presence of those mechanisms in a particular organization.

Representational standing can be analyzed along four dimensions. *Access* concerns the practical opportunity to submit evidence. *Intelligibility* concerns the availability of interpretive categories and linguistic support through which that evidence can be understood. *Credibility assessment* concerns the grounds and procedures used to appraise a claim. *Consequential recognition* concerns whether the result can affect the institution’s records, inquiries, decisions, or reasons for action. These dimensions can diverge. An intake form may be accessible while its taxonomy cannot express the submitted concern; testimony may be recorded but assigned negligible weight; or a claim may be accepted as relevant without reaching a competent decision-maker. The proposed dimensions describe an audit interface and do not constitute a validated measurement scale.

6.2 Classification authority and interpretable claims

Classification systems define practical boundaries around the events, persons, and relationships that institutions can register. Bowker and Star describe the social consequences of classifications and standards, including the work by which categories become routine and their assumptions become less visible [7]. A representation inherits some of these choices whenever its variables and labels are constructed from administrative records. The effect of a category depends on the institutional consequences attached to it:

eligibility, review competence, allocation, or public visibility may turn on a coding decision. The same label may carry different consequences across institutions and historical configurations.

Consider a research consortium whose reporting system recognizes funded investigators, contracted technicians, and formally registered community contributors. A person whose practical contributions fall outside these categories may be included in an aggregate work-hour figure while remaining ineligible for an authorship inquiry or a dispute procedure. The numerical record may be correct relative to the recorded hours. The missing distinction concerns the institutional effect of classification on an affected subject's ability to advance a claim. A defensible analysis would examine the actual terms of contribution, applicable rules, the person or body's competence to interpret them, and any alternative procedure. Classification disparity alone does not establish misconduct or an enforceable entitlement.

A category can also offer protection. Requirements to distinguish sources, mark evidential uncertainty, separate confirmed and alleged events, or conceal sensitive identities can improve both reliability and safety. Accordingly, the ethical inquiry is directed toward the authority, purposes, evidential warrant, contestability, and consequences of a classification. The question is which distinctions an institution needs to preserve for the decisions it is authorized to make, and whose knowledge can influence the specification of those distinctions. D'Ignazio and Klein's analysis of power in the production and use of data supplies an intellectual precedent for this concern [9]. The present work adds a narrower interface for connecting classification choices to the evaluability of declared institutional questions.

6.3 A staged model of representational admission

For a declared research purpose, let D denote a set of admissible candidate histories, each characterized by explicitly stated evidential and interpretive conventions. Let π denote a representational policy covering an intake procedure, a classification practice, an evidential evaluation rule, and a reporting design. The resulting observation is modeled by the composition in Equation 10.

$$\Phi_\pi : D \longrightarrow Z_\pi, \quad \Phi_\pi = G_\pi \circ W_\pi \circ C_\pi \circ E_\pi. \quad (10)$$

Here $E_\pi : D \rightarrow \mathcal{E}_\pi$ denotes evidence accessible through a specified intake and collection arrangement, $C_\pi : \mathcal{E}_\pi \rightarrow \mathcal{C}_\pi$ denotes a declared coding or interpretive procedure, $W_\pi : \mathcal{C}_\pi \rightarrow \mathcal{R}_\pi$ denotes an evidential appraisal producing an institutional working record, and $G_\pi : \mathcal{R}_\pi \rightarrow Z_\pi$ denotes the public or decision-facing reduction. The mapping is deterministic for the limited formal argument below. Actual testimony is interactive, credibility assessments are often contestable, and record-production mechanisms may be stochastic or institutionally distributed. Such cases require richer representations. The functions do not imply that a person's experience can be reconstructed exhaustively from records.

A proposition about the final reporting map concerns Φ_π on its declared domain. It is consistent with a failure of the earlier intake or interpretation procedures. Conversely, the observation of a difference in a published outcome does not determine whether a subject was excluded at intake, misclassified, disbelieved, or affected by a lawful and appropriately reasoned eligibility rule. The same aggregate pattern may be compatible with materially different mechanisms. Investigating these mechanisms therefore calls for records of the stages themselves, interviews or other suitable evidence, and competent review of the relevant institutional rules.

For a specific question q , let $J_q : D \rightarrow A_q$ denote a declared target interpretation, such as whether a particular form of standing or access was available in a stipulated history. An audit confined to the published output can recover the answer exactly only when the target is constant on the attained fibers of Φ_π . This elementary criterion is used in Section 6 to locate a concrete epistemic limit; it is inherited from ordinary function factorization and from the preservation analysis in Section 5.

6.4 A limit of output-only auditing

A common audit strategy examines the published representation and tests whether its totals, comparisons, or classifications are internally consistent. Such a test can identify consequential errors within

the available data. Its scope is narrower when the suspected mechanism concerns who could provide the underlying evidence or how testimony was received. Proposition 6.1 makes this limit explicit under simple assumptions.

Proposition 6.1 (Output-only recovery of a declared claim). *Fix a deterministic representational policy π , a domain D , and a target question q with answer map $J_q : D \rightarrow A_q$. If there exist $h_1, h_2 \in D$ with $\Phi_\pi(h_1) = \Phi_\pi(h_2)$ and $J_q(h_1) \neq J_q(h_2)$, then no function of $\Phi_\pi(h)$ alone returns $J_q(h)$ correctly for every $h \in D$.*

Proof. If a function $a : Z_\pi \rightarrow A_q$ recovered the target on all histories, it would assign the same value to $\Phi_\pi(h_1)$ and $\Phi_\pi(h_2)$. Its correctness would require those identical outputs to equal both distinct target values, a contradiction. $\square$

Proposition 6.1 identifies a limit on available information, without diagnosing the social mechanism producing the equivalence. Its conclusion is conditional on the stated domain, mappings, and interpretation of J_q . The proposition does not demonstrate identity prejudice, testimonial silencing, intentional concealment, or an institutional obligation to disclose confidential information. A reviewer with legitimate access to additional intake documents or protected testimony may distinguish the two histories. A research design can therefore compare the conclusions supported by public summaries, authorized internal records, and independently collected evidence.

The epistemic significance of the formal result lies in the location of the audit’s evidential dependence. Adding statistical precision to an output may improve estimates of the variables it contains while leaving missing avenues of testimony untouched. Likewise, retaining source identifiers may improve traceability while the reasons for refusing a claim remain inaccessible. A claim of epistemic injustice requires a further account of credibility, interpretive disadvantage, social power, and the circumstances of the particular institutional interaction. The formal result disciplines assertions of recoverability; the normative diagnosis depends on independent reasons.

6.5 Constructed comparison of complaint interpretations

Consider a hypothetical environmental-monitoring office responsible for responding to reports about water quality. Its public summary records the number of complaints received, the number marked complete, and the median time to administrative closure. The summary is compiled from a protected case register. The following cases are deliberately constructed; they describe no identifiable public body, affected group, or actual environmental occurrence.

In history h_A , residents submit observations in a form recognized by the official taxonomy. The office records an inspection, communicates uncertainty about its findings, and offers a defined review pathway. In history h_B , some residents use an account of environmental change for which the intake interface provides no suitable category. Their submissions are filed under general correspondence and excluded from the complaint denominator. In history h_C , their observations reach the complaint register but the reviewing office treats them as insufficiently credible without documenting the criteria used. A fourth history h_D contains the same public totals but a documented and justified competence limitation, alongside a referral that preserves meaningful access to another review body.

The histories can be constructed so that the published complaint totals are equal. They differ in the mechanism of recognition and the availability of consequential institutional inquiry. A researcher assessing h_A through h_D would need to determine whether the concerns are evidentially supported, whether official categories were reasonably available, whether a credibility judgment followed appropriate procedures, and whether an alternative competent pathway was effective. The example invites investigation of representational mechanisms while deliberately withholding a categorical finding of injustice.

Table 6 separates candidate mechanisms from discriminating observations. The table is a proposed research instrument whose entries require adaptation to actual institutional roles, records, and applicable confidentiality requirements.

The table also illustrates the difference between the conceptual diagnosis and the identification of an actual cause. For example, a discrepancy between intake forms and public counts may result from

Table 6: Possible mechanisms of representational exclusion and evidence for distinguishing them.

Mechanism	Observable difference sought	Limits or rival explanation
Practical intake barrier	Unsubmitted accounts, failed access attempts, or inaccessible language channels	Differential awareness, capacity, and voluntary nonparticipation
Interpretive mismatch	Repeated use of residual categories or claims lacking a recognized description	Ambiguous evidence or categories adequate for the authorized task
Credibility disparity	Different evidential treatment of comparable testimony and reasons offered for it	Legitimate differences in corroboration, competence, or reliability
Administrative reclassification	Change in eligibility or procedural routing associated with a coding decision	Lawful jurisdictional rule with functioning referral safeguards
Aggregation loss	Public equivalence of cases with distinct standing or remedy paths	Protected records or independent review may preserve the distinction
Remedial exclusion	Corrected records without a competent route to consequential reconsideration	Valid jurisdictional limits or a separate effective appeal channel

delayed synchronization, sampling, or coding changes. Establishing unfair exclusion would demand more than a discrepancy. The institutional record should be interpreted with evidence of the claimant’s position, the office’s competence, applicable standards of proof, reasons for classification, and alternative procedures. Testimonial evidence can be indispensable to that investigation without enjoying automatic immunity from scrutiny.

6.6 Disclosure, interpretation, and protected testimony

Sensitive testimony presents a further difficulty. Publicly identifying a claimant may create exposure to retaliation, stigma, or unintended circulation of private information. The ethical goal of increased visibility must be evaluated alongside confidentiality, autonomy, and context-specific expectations governing information use. Nissenbaum’s account of contextual information norms provides a relevant framework for assessing appropriate circulation [33]. In many settings, a public count, a protected case record, and an independently reviewable decision can serve different legitimate purposes.

Three protections deserve separate assessment. First, a protected submission pathway may permit testimony in forms adapted to the affected population, including interpretation or accessible formats. Second, an authorized reviewer may inspect confidential evidence and assess the reasons for inclusion, exclusion, or credibility judgments. Third, a public-facing explanation may disclose the procedure, aggregate disposition, and limits of the record without exposing personal details. These measures do not guarantee a remedy: their significance depends on competence, independence, reasonable access, and a genuine ability to change a consequential decision. Those issues form part of the representational-contestability inquiry developed in the next section.

Provenance metadata can identify stages of record production, transformations, and the responsible institutional units, using vocabularies such as PROV-DM or PROV-O [46, 47]. Such metadata can support procedural inquiry. Its usefulness depends on whether the recorded provenance includes contested exclusions and whether authorized reviewers can inspect the relevant details. Complete provenance of accepted records is compatible with inadequate observation of accounts that never entered the record. A responsible provenance system may therefore document admission rules, coding versions, reasons for revisions, and the limits of what could be observed, while avoiding unsupported claims to comprehensive coverage.

6.7 Differential evidence for silence, withholding, and non-observation

The mechanisms of representational exclusion call for a more discriminating procedure when a public record contains no trace of a potential claim. The objective of the present subsection is to distinguish patterns of institutional non-observation from practices that restrict a claimant's epistemic standing. Its method is a comparative evidential assessment: competing accounts of a missing claim are specified, possible discriminating observations are identified, and the permissible means of obtaining those observations are assessed independently. The resulting procedure remains a research proposal. It provides neither a diagnosis of injustice from an empty field nor a general entitlement to collect additional personal information.

A missing entry has at least three analytically distinct referents. First, the underlying account may never have been presented to the relevant institution. Second, an account may have been received but become unrecognizable or ineligible through interpretation, coding, or assessment. Third, the account may be retained within a protected channel and appropriately absent from the public representation. Each possibility admits further differences. Non-submission may reflect a meaningful opportunity to decline participation, a reasonably perceived risk of retaliation, unawareness of the available procedure, or practical inaccessibility. A protected record may result from appropriate confidentiality measures or from a contested restriction on access to review. The same observable public absence is consequently consistent with contrasting ethical assessments.

Fricker's distinction between testimonial and hermeneutical injustice is especially relevant to the second possibility: the conditions for an epistemic wrong concern credibility prejudice or unequal interpretive resources, with contextual qualifications that exceed the mere absence of a category [16]. Dotson's investigation of testimonial silencing directs attention to the interaction between a speaker's vulnerability and a hearer's response, including practices that cannot be diagnosed by examining a published dataset alone [11]. Medina's work on epistemic resistance likewise emphasizes the conditions under which differently situated subjects can participate in practices of knowing [30]. These analyses motivate a search for relevant encounters and interpretive practices; they supply no presumption that an absent submission proves wrongful silencing.

Table 7 proposes a differential evidence register for a specified institutional question. The observations are potential indicators to be interpreted with the affected parties and applicable procedural rules, rather than a standardized set of diagnostic tests. Each comparison should record whether a proposed observation is accessible through ordinary public documentation, requires authorization for protected records, or would depend on voluntary participation.

The register suggests a staged audit under explicit safeguards. A researcher can first delimit the institution's actual mandate, the question for which the representation is used, and the population whose claims may be consequential. The next stage compares accessible descriptions of intake channels, classification schemes, and aggregate outputs, recording gaps and variations without attributing motives. Where a consequential discrepancy remains, an independent or otherwise competent reviewer may examine access-controlled logs and reasons for coding or exclusion. Separately, affected subjects may be offered confidential and genuinely voluntary opportunities to describe how the intake and interpretation arrangements operated. Such opportunities must include nonparticipation and withdrawal options appropriate to the study, with an assessment of power differences, translation burdens, and risks of retaliation. The researcher then compares the resulting explanations against documentary evidence, procedural rules, and rival accounts. Each stage may produce an indeterminate result; an indeterminate assessment can be the most defensible finding when records cannot safely be obtained.

The resulting conclusions require graduated language. Evidence of an inaccessible form supports a claim about the form's availability in the observed context. Evidence that a received submission was assigned a residual code supports an account of coding practice. A finding of testimonial injustice additionally requires an argued connection between a credibility deficit, prejudice, and the person's epistemic standing. An allegation of silencing calls for examination of the relevant testimonial interaction and vulnerabilities. A claim of institutional wrongdoing requires identification of the applicable obligations, competence, and remedies. None of these inferential steps is secured by statistical disagreement between

Table 7: Differential evidence register for missing or inaccessible institutional claims.

Candidate account	Discriminating evidence sought	Interpretive and access limitation
Practical non-access	Availability and usability of intake channels; documented unsuccessful attempts	Barriers may coexist with limited awareness, resource constraints, or unrelated nonparticipation
Credibility-dependent exclusion	Comparable submissions, appraisal criteria, reasons and review outcomes	Different evidential strength or jurisdiction can warrant different treatment
Interpretive mismatch	Coding manuals, residual categories, contested translations and requested recoding	The declared task may reasonably delimit categories; adequacy is question-relative
Voluntary withholding	Consent terms, independently accessible channels, noncoercive opportunities to participate	Silence should not be treated as consent; obtaining evidence cannot require compelled disclosure
Protected retention	Restricted-case register, authorized review procedures and publication policy	A public omission can be appropriate; protected records cannot be assumed accurate or complete
Administrative or technical non-observation	Version history, synchronization logs, retention schedules and data-source coverage	A missing record may arise from ordinary processing error without determining a moral diagnosis

public totals. Even consistent testimony about an exclusion mechanism needs contextual corroboration where available, without imposing an unrealistic demand that vulnerable participants repeatedly disclose private experiences.

The protective dimension is substantive. Nissenbaum’s contextual-integrity account evaluates information flows in relation to contextual norms, actors, information types and transmission expectations [33]. For the proposed audit, the legitimate interest in making excluded claims examinable must therefore be distinguished from a demand for general publication. An authorized evaluator may verify that protected submissions were considered while publishing only non-identifying descriptions of the procedure and its limitations. The evaluator should also document what cannot responsibly be inferred from redacted material. An audit that increases personal exposure or makes participation a condition of access to remedy could compound the very disadvantage it seeks to investigate. D’Ignazio and Klein’s analysis of power in data practice reinforces the need to examine whose effort and risk a new evidence requirement imposes [9].

Consider two constructed reporting offices that publish identical counts of unresolved complaints. One maintains a confidential intake channel requested by contributors who wish to avoid public identification; authorized reviewers can inspect its restricted records. The other accepts submissions only through a public form that certain affected subjects reasonably consider unsafe, and offers no comparably accessible protected route. The published counts alone are insufficient to distinguish these arrangements. A protected review of intake rules, refusal reports, authorized records and freely offered participant accounts could support a more informative comparison. The example also admits rival interpretations: the apparent safety concern might reflect inadequate communication rather than a design defect; an existing alternative channel might be effective; or the evaluator might lack access to decisive evidence. Thus, the case remains deliberately hypothetical and causally indeterminate without additional observations.

Four classes of counterevidence would bear directly on a proposed exclusion diagnosis. A verified, usable protected channel may weaken a claim that public exposure was the only means of participation. Consistently documented evidential criteria may weaken an allegation of discriminatory credibility assessment, provided their formulation and application withstand independent scrutiny. A review mechanism that reaches competent decision-makers may rebut a narrower claim of total remedial inaccessibility. Evidence that renewed data collection would expose claimants to unreasonable risk may defeat the appropriateness of a proposed public transparency measure. These findings can restrict one hypothesis

while leaving other ethical questions open.

The proposal's incremental contribution is procedural: it combines differential explanations of silence with explicit constraints on permissible evidence acquisition and with bounded conclusions about institutional authority. Its utility should be compared with competent qualitative inquiry, independent grievance review, and existing documentation practices. Where those methods already establish the relevant distinction safely and at reasonable cost, the present register may add little. Its central value is the discipline of treating an absence as a question requiring contextual investigation, while recognizing that some information may properly remain inaccessible to the researcher or the public.

6.8 Competing explanations and responsible interpretation

A defensible research method requires alternatives to an epistemic-injustice explanation. Differences in response rates can arise from awareness, geography, access costs, or selection into the studied population. A credibility assessment can be supported by corroboration standards relevant to the decision. A missing category can reflect a provisional data model whose declared purpose is limited. An institution may also restrict publication to protect participants while retaining effective review. These possibilities neither eliminate the risk of exclusion nor establish that a reported disparity is ethically acceptable. They specify the factual and normative questions that a credible diagnosis must answer.

At least four observations could weaken a particular representational-exclusion hypothesis. Records may show that allegedly excluded submissions were evaluated through a competent alternative channel. Protected review may establish that disputed classifications rested on consistently applied and relevant evidential criteria. An independently accessible account of the affected subjects' experiences may contradict the researcher's proposed interpretation of the missing distinction. A comparison of admissible data designs may show that the proposed refinement would not improve the targeted decision or would create disproportionate confidentiality risks. Each finding warrants revising a specific inference, even where broader institutional concerns persist.

The competing methods include qualitative interviewing, participatory inquiry, archival reconstruction, conventional administrative audit, and legal evaluation. A staged representational model can organize information about where a loss or institutional decision occurred. The model's value must be judged by whether it clarifies causal or normative disagreements, supports scrutiny of errors and exclusions, and improves decision-relevant evidence at reasonable cost. It has no priority over methods that recover the same distinction more directly and responsibly. The comparison should also consider the labor of participants who supply missing evidence, the risk of demanding repeated disclosure, and the possibility that formal categories reproduce the very asymmetries they are designed to make visible.

6.9 Analytical limits and institutional implications

The section's argument establishes an interface between epistemic injustice and representational practice. Institutions select methods by which accounts enter records, acquire classifications and credibility judgments, and become eligible for aggregation or decision. The formal pathway makes those choices explicit, while Proposition 6.1 identifies a limited information boundary for audits confined to output. Neither supplies an empirical finding of wrongdoing or a general criterion of legitimate testimony. A normative assessment requires attention to rights, authority, interpretation, reliability, possible harms, and the historical conditions of participation.

The proposed practical implication is a duty of inquiry whose content is contingent on the purpose and foreseeable consequences of the representational system. Where representations inform consequential decisions, a researcher or institution has reason to examine what evidence the model admits, which claims may remain outside its recognized categories, how credibility is assessed, and how protected review can address documented failures. This is a proposed ethical orientation, with its force and institutional implementation requiring case-sensitive argument rather than derivation from information theory.

Section 6 has concentrated on the mechanisms by which a representation acquires its record of claims. Section 7 addresses the allocation of power to challenge those decisions, obtain a reasoned response,

revise classifications or admissibility rules, and track whether revisions reach consequential uses of a model. That transition retains the distinction between an epistemically detectable omission and a justified institutional remedy.

7 Contestability of Representational Decisions

This section examines the institutional conditions under which an affected subject can challenge the construction, interpretation, and consequential use of a reduced representation. Its objective is to identify distinct objects of contestation, trace the authority and evidence required for corrective review, and specify the relationship between record revision and the reconsideration of decisions that relied on an earlier version. The discussion first separates technical correction from procedural and substantive remedies and examines conditions for adapting contestation procedures across institutional mandates; it then proposes a version-sensitive model of distributed representational use, establishes an elementary propagation result, and evaluates a constructed institutional comparison. It concludes with safeguards, discriminating evidence, and limits on the transfer of procedures between settings. The models are analytical devices and do not by themselves create rights, establish institutional jurisdiction, or determine ethical legitimacy.

7.1 Objects and standing of representational challenges

An institutional representation is contestable along several dimensions, each associated with a distinguishable evidential question. A subject might challenge whether their account was received; how the account was classified; which records were admitted to a summary; whether a descriptive reduction preserves an important relation; or whether a decision-maker was justified in relying on that summary. Challenges to the source record and challenges to its authorized uses can therefore concern different decision-makers and require different kinds of evidence. A corrected administrative label does not, by itself, establish the validity of a benefit claim or alter a prior allocation. Conversely, a competent body may decide that an adverse allocation should be reconsidered while leaving a faithfully recorded historical statement intact.

The mechanisms analyzed in Section 6 suggest why the availability of a complaint form provides only limited evidence of meaningful contestation. A person who lacks an admissible category may also lack formal standing to request the correction of that category; a claim may be received by the data steward but remain outside the competence of the authority responsible for the consequential decision. Bowker and Star’s work on the institutional effects of classification, and the analysis of abstraction in sociotechnical systems developed by Selbst and colleagues, provide precedents for examining these distributed dependencies [7, 43]. Their work supports the importance of contextual assessment. It does not entail that every categorization confers standing or that every claimant possesses an identical remedy.

For present purposes, *representational contestability* denotes the demonstrable possibility, within declared institutional rules, of submitting a reasoned challenge to a relevant representational decision, obtaining its consideration by an authorized reviewer, and receiving an intelligible disposition. Its effectiveness depends on competence, accessibility, the ability to consider relevant evidence, and the consequences that the authorized decision can produce. An institution may legitimately impose evidential standards, deadlines, and confidentiality controls. Their ethical assessment depends on the interests at stake, burdens they impose, alternatives available, and applicable legal constraints. Claimants, witnesses, data stewards, domain experts, and downstream decision-makers need not occupy interchangeable roles in the process.

The legal and governance literature offers important, context-dependent precedents. Kaminski and Urban develop four contestation archetypes by distinguishing rules from standards and procedural from substantive emphases; their analysis proposes and compares arrangements for challenging algorithmic decisions [24]. Kroll and colleagues discuss mechanisms for verifying compliance of automated decisions with governing standards [26]. Both analyses primarily address computational decision-making, whereas the present paper also examines manually constructed institutional summaries and interpretive

classifications. The OECD’s due-diligence guidance separately emphasizes legitimate grievance mechanisms, their scope, accessibility, implementation, and monitoring in responsible-business-conduct settings [34]. The present framework draws a methodological distinction across these literatures; it does not transfer a particular legal obligation to every representational practice.

7.2 Authorization, evidence, and review competence

A representational challenge can be described through a documentary *review warrant*. The warrant records the challenged representation, the requested change, the evidence offered, the reviewing body’s competence, the reasons supporting the disposition, and any recognized pathway to further inquiry. This is a proposed audit instrument rather than a newly established legal doctrine. Its purpose is to permit an investigator to distinguish the existence of a review procedure from its capacity to affect institutional knowledge and action.

A useful audit begins by identifying the object of authority. A records custodian may change transcription errors; a technical maintainer may alter an aggregation algorithm; a professional committee may determine eligibility within its mandate; and a court or administrative tribunal may possess authority to review a legally consequential decision. The relationship between these roles must be established for the institution under examination. A procedurally authorized correction of a dataset does not necessarily confer the power to amend a separate decision. Equally, a downstream authority may have reasons to review a decision on grounds that do not require rewriting the underlying record.

Evidence should be assessed at the level appropriate to the challenge. A disagreement over a recorded event may require source documents and witness accounts. A challenge to a classification may require the relevant taxonomy, versions of coding instructions, and records of applied exceptions. A dispute over a mathematical reduction may be resolved by demonstrating that the mapping loses a declared distinction, as in Section 4. A challenge to a normative rule or the competence of a decision-maker requires legal, institutional, and ethical argument extending beyond the mathematics of that reduction. Multiple grounds can coexist; a record should therefore disclose which issue a given reviewer actually decided.

The European Union’s General Data Protection Regulation provides a specific example of legally defined protections for certain solely automated individual decisions with legal or similarly significant effects. Article 22(1) specifies the covered class, subject to the exceptions in Article 22(2). Article 22(3) explicitly requires safeguards, including human intervention, an opportunity to express a view, and contestation, in the cases falling within Article 22(2)(a) and (c) [14]. The provision is jurisdiction- and subject-matter-specific, and the cited text does not establish a general right to correct every normative representation or reopen every adverse institutional decision. Its relevance here is architectural: meaningful review must be situated within an applicable legal or institutional mandate rather than inferred from the mere existence of an explanatory interface.

7.3 Procedural Portability and Competence Boundaries

Institutional designers sometimes borrow the language of contestation from a neighboring legal or technical regime. A source-specific comparison is needed before an existing procedure can serve as a model for a new representational setting. This subsection examines the transfer of procedural design principles across automated decisions, voluntary AI-governance frameworks, and manually assembled institutional records. Its objective is to specify the premises required for a transfer, while distinguishing the evidential accessibility of a procedure from the scope of an enforceable entitlement.

Three source domains expose different features of procedural design. Article 22 of the GDPR addresses decisions based solely on automated processing that produce legal or similarly significant effects, subject to its specified exceptions; Article 22(3) identifies human intervention, the opportunity to express a view, and contestation for the cases covered by Article 22(2)(a) and (c) [14]. Kaminski and Urban develop rule-based, standard-based, procedural, and substantive dimensions of contestation in considering algorithmic decisions; the resulting archetypes offer analytical choices without specifying the authority of every potential reviewer [24]. The OECD AI Recommendation, amended in May 2024, encourages

Table 8: Scope checks when adapting procedures for challenges to consequential representations.

Precedent and scope	Potentially transferable design feature	Required contextual verification
GDPR Article 22: specified solely automated decisions	Human intervention and opportunity to contest under the stated statutory conditions	Applicable jurisdiction, legal effects, paragraph 2 exception, paragraph 3 safeguards and competent body
Kaminski–Urban: algorithmic contestation scholarship	Alternative procedural and substantive review designs	Authority, standing, practical access, effectiveness and scope beyond algorithms
OECD AI Recommendation: AI outputs	Context-sensitive disclosure facilitating challenge	Adoption into the relevant practice and any applicable enforceable basis
NIST AI RMF 1.0: voluntary AI governance	Lifecycle documentation, monitoring and organizational risk review	Institutional adoption, permitted uses and authority for consequential reconsideration
Manually constructed institutional records: proposed application	Claim intake, versioned correction, protected review and downstream notification	Identifiable competent reviewer, legal constraints, actual reliance and availability of remedy

context-appropriate information enabling those adversely affected by AI systems to challenge outputs; as a recommendation, it leaves the creation and enforcement of particular legal entitlements to applicable arrangements [36]. These sources contribute different kinds of guidance, and a comparison should preserve that heterogeneity.

A transfer analysis should begin with a *trigger*: whether the challenged act involves a regulated automated decision, an internal classification, a public aggregate, or a downstream determination using that aggregate. The investigator should then identify the *subject and standing rule*, the *competent reviewing body*, the *evidence that body may lawfully examine*, and the *kind of disposition it may issue*. A further question concerns the relationship of that disposition to a second authority using an earlier representation. In a multi-body system, a data steward may have competence to amend a record, while a statutory decision-maker retains discretion or a duty governed by a separate procedure. The issue can be studied empirically through records of competence, notification, reason-giving, and actual version use.

Table 8 organizes these tests. Its rows are sources of potentially instructive procedural architecture, rather than declarations that each source governs all institutional models. In particular, the analogy with the voluntary NIST AI Risk Management Framework is relevant to documenting model governance and monitoring, subject to its stated AI-system scope [31]. The classification of a procedure as “contestable” is incomplete until the legally or organizationally competent reviewer and available disposition have been established.

Consider a hypothetical municipal eligibility list created through staff coding of paper submissions. The list may influence benefit decisions without itself constituting a solely automated decision. Incorporating an understandable explanation and a channel to challenge a coding error might be institutionally worthwhile, even where Article 22 does not govern that record. A finding that the list contains a recoverable classification error establishes a technical question. To establish an obligation to reopen a benefits decision, the analysis additionally requires the applicable procedural rule, the decision-maker’s competence, and evidence about whether the corrected information was relevant. By contrast, an existing independent appeal based on original submissions may render a separate version-propagation process unnecessary for an effective review. Both cases are relevant to assessing the incremental contribution of the present framework.

The appropriate research output is therefore a *scope-qualified procedural claim*: an identified object, evidential basis, competent authority, authorized disposition, and an explicit statement of what remains unsettled. Legal scope, organizational practice, and ethical assessment each enter the inquiry with distinguishable warrants. This bounded approach permits transfer of design ideas while reducing the risk that a philosophical account of contestability is mistaken for a declaration of general legal rights.

7.4 Versioned corrections and consequential use

A representation often circulates beyond the body that originally compiled it. Published figures may be copied into reports, translated into eligibility classifications, supplied to oversight bodies, or embedded in subsequent allocation procedures. Correcting the source record creates a new evidential state, but the practical effects depend on which versions the receiving institutions use and whether they possess authority and procedures to reconsider their own outputs. A careful description distinguishes *record correction*, *revision of a representational policy*, *notification to users*, and *consequential reconsideration*.

Let H be a declared set of candidate histories, and let $R_{\pi,v} : H \rightarrow Z_{\pi,v}$ describe a representation obtained using policy π and an identified version v . A correction request e may lead to a revised policy and version, represented by $U(\pi, v, e) = (\pi', v')$, where the operation is defined only for accepted and authorized requests. For each declared downstream use $j \in \mathcal{J}$, let σ_j identify the version actually available to that use, c_j its relevant local information, and G_j its stipulated decision rule. Equation 11 separates the corrected record from the version on which a downstream output depends.

$$z' = R_{\pi',v'}(h), \quad d_j = G_j(R_{\sigma_j}(h), c_j), \quad j \in \mathcal{J}. \quad (11)$$

Here the abbreviated R_{σ_j} includes both the applicable policy and the selected version. A downstream body may use a superseded but still legally operative version, or may have adopted the corrected representation without reopening a completed decision. Such possibilities cannot be resolved from the corrected source alone. In actual settings, changes in evidence, legal rules, and review competence may also modify G_j or c_j ; Equation 11 holds these separate so that a particular mechanism can be examined.

A *consequential correction trace* therefore records the accepted change, the reasons and competent authority supporting it, the material uses reasonably implicated, notification or propagation events, and the dispositions of any subsequent review. The analytic endpoint is appropriately specified by the institutional question. Some circumstances require accurate republication; others require reconsideration of a decision or an explanation for declining to reopen it. None of these outcomes necessarily requires that every downstream decision change its substantive result. A reasoned finding that a correction does not alter the applicable decision may represent a meaningful review, provided the reviewer has competence and has considered the corrected evidence.

7.5 Dependency reachability and limits of correction

It is useful to formulate a narrow structural requirement for the propagation of corrections. Let $\mathcal{G} = (V, E)$ be a finite directed acyclic graph of declared representational dependencies for a particular audit. Each vertex denotes a *record-producing or decision-producing operation*; edges identify an authorized route by which an operation's output is used as an input elsewhere. This graph is a description of dependence among operations. Its direction does not assert a universal temporal order, an ontological causal structure, or moral priority. The model stipulates fixed local inputs and deterministic operations for the limited result below.

Proposition 7.1 (Reachability required for a propagated correction). *Consider two evaluations of the dependency graph $\mathcal{G}$ that differ only in the output supplied at a designated vertex s . If a vertex t is unreachable from s by a directed path, then its computed output is identical in the two evaluations. Hence a correction at s can alter the computed output at t only if the active dependency graph contains a directed path from s to t .*

Proof. Because $\mathcal{G}$ is finite and acyclic, its operations can be evaluated in a topological ordering of their dependencies. Every vertex outside the descendants of s has unchanged local inputs and receives inputs from vertices that are also outside those descendants. Starting with input vertices and proceeding by induction through such an ordering, its computed output is unchanged. The conclusion follows for t . The proof applies to the stipulated deterministic dependency model and gives a necessary condition only. $\square$

Proposition 7.1 establishes an elementary graph-dependency result. Reachability is insufficient for a substantive outcome change: the receiving operation may ignore the corrected field, apply a threshold insensitive to the revision, or have a competent reason for leaving its decision intact. Nor does a graphical path establish that a body is legally obligated to reopen a decision. Conversely, where a declared use relies on a separate information channel, the purported dependency graph may be incomplete. The investigator must then revise the graph before inferring anything from the proposition. Cyclic updating procedures, discretionary decisions, and stochastic records require additional assumptions and are outside its simple proof.

The model offers a modest but useful empirical diagnostic. Investigators can compare a declared dependency register with version logs, notifications, access records, and the actual files inspected by the decision-maker. They can then distinguish an absent correction pathway from an ineffective pathway and from a pathway that functioned but led to a justified unchanged outcome. A failure to propagate a corrected representation becomes ethically consequential when it interacts with a relevant claim, a duty of reconsideration, or foreseeable harms. These relations require independent evidential and normative argument.

7.6 A constructed case of correction and institutional reliance

Consider a hypothetical regional environmental monitoring consortium that produces a public assessment used by a separate fund to allocate restoration grants. The consortium records site reports, community submissions, measurement protocols, and geographic classifications. The fund uses a simplified risk indicator to rank eligible projects under a declared programme. Three institutions participate: a data steward, a scientific review panel, and a grant-allocation committee. Their procedures and authority are stipulated for the example; the case refers to no actual organization or persons.

A community group submits evidence that one monitoring location was assigned to the wrong catchment zone. The data steward accepts the correction and issues a new dataset version. In history h_1 , the review panel records the change but the fund's committee continues to use an archived risk table, without having been notified of the new version. In history h_2 , the corrected dataset reaches the fund and triggers reconsideration, although the fund's ranking remains unchanged because other eligibility criteria are controlling. In history h_3 , the updated risk indicator changes the ranking, but the committee has no authority to revise allocations already fixed under the relevant programme. A fourth history h_4 uses a protected challenge pathway, records the committee's competence decision, and gives the claimant an intelligible referral for any further review authorized by the programme.

All four histories can share the same publicly displayed corrected catchment label. They can nevertheless differ in notice, competent review, eligibility, and implemented relief. The example shows why the truth of a correction, the reviewability of an allocation, and the eventual provision of a remedy should be recorded separately. It also illustrates a negative case: even an apparently significant representational change can leave a responsible decision unchanged after effective review. The ethical question concerns the actual reasons, jurisdiction, burdens and available alternatives, not a presumption that every successful contest must alter the outcome.

Table 9 identifies discriminating evidence for these possible mechanisms. The table is proposed as an aid to qualitative and documentary inquiry. It is neither a scoring instrument nor an exhaustive list of legal requirements.

7.7 Protected review and unequal burdens of contestation

Representational challenges can impose significant burdens on affected subjects. A formal opportunity to contest can require specialized literacy, access to records, time, legal assistance, or repeated disclosure of sensitive experiences. The information required to substantiate a challenge may be controlled by the very organization whose classification is disputed. A credible assessment therefore examines practical access and the distribution of evidential costs. Consultation with affected people is relevant to the design

Table 9: A proposed audit of representational challenge, correction, and downstream consequence.

Stage of challenge	Evidence needed for evaluation	Alternative interpretation or limit
Access and standing	Accessible submission channel, eligibility rules, language and confidentiality provisions	Competence may lie with another body offering a functioning referral
Evidence and classification	Versions of codebook, source submissions, corroboration criteria, written reasons	Correction may address a transcription error without changing entitlement
Authorized disposition	Reviewer mandate, reasoned finding, correction authorization, appeal options	Reviewer may be competent for records but not downstream allocation
Propagation of versions	Version identifiers, recipient registers, notification logs, actual files used	Recipients may use an earlier version pursuant to a lawful rule
Consequential reconsideration	Reopened decision records, evidence considered, reasons and implementation	Corrected evidence may lead to a reasoned unchanged disposition
Longer-term accountability	Record of unresolved issues, independent scrutiny, evidence of realized effects	Monitoring alone cannot establish a causal or ethical improvement

of procedures, but consultation alone does not establish representation, consent, or a transfer of decision authority.

A protected review architecture may separate access to confidential records from publication of aggregate findings. The authority receiving sensitive testimony should have a defined mandate, security arrangements, and a route to appropriate corrective action. Public explanations can identify the challenged policy, procedural safeguards, and reasons for disposition at a level consistent with legitimate confidentiality constraints. Nissenbaum’s work on contextual information norms helps explain why audiences and uses matter to the permissibility of disclosure [33]. The W3C provenance specifications describe how records of activities, entities and their relationships may be represented [46, 47]; their vocabularies can support traceability without resolving disputes over authority or the truth of an evidential claim.

Several institutional design tensions remain. Publication of a detailed reason may disclose a protected claimant; suppression of the reason may undermine the ability to challenge an outcome. Broad appeals may improve access but create delay, cost, or strategic use of procedures. A single central reviewer may offer consistency while remaining distant from the relevant community’s interpretive resources. These are matters for institutional design and situated ethical assessment. Suitable safeguards may include structured reason-giving, access to independent or appropriately separated review, confidentiality-preserving evidence examination, staged correction, and documented routes for reopening decisions where authorized. Their value depends on implementation and the circumstances of the claim.

The possibility of deliberately adversarial submissions should also be considered. An institution can justifiably require relevant evidence, protect participants from retaliatory exposure, and prevent harassment of reviewers. These safeguards should themselves be open to proportionate scrutiny. Otherwise, security or evidential rules may become a mechanism by which affected subjects are denied meaningful consideration. The empirical task is to examine whether protective constraints are applied consistently and whether competent alternatives permit the relevant claim to be heard.

7.8 Evaluation, rival explanations, and institutional boundaries

An institutional comparison should distinguish the registration of a challenge, acceptance of a correction, competent reconsideration, substantive revision of a decision, and implementation of any remedy. These events may be ordered differently in separate institutional procedures and may involve multiple authorities. Their documentary sequence is not itself proof that one caused another. A revised funding

outcome may arise from changing budget conditions, new scientific evidence, external enforcement, or discretionary prioritization. Assessing the effect of contestability calls for an explicit contrast between these explanations and the mechanisms the model identifies.

A proposed comparative design would specify the relevant claimant population, the authority of each reviewing body, the versions of information actually used, the decision rules, and the circumstances under which prior decisions can be reopened. Evidence should include unsuccessful, withdrawn, and unregistered complaints where access can be investigated responsibly. Researchers could compare similar challenges routed through different procedures, examine corrections that had no downstream consequence, and assess the reasons for decisions left unchanged. Qualitative interviews, process tracing, legal analysis, and independent technical audit are complementary methods. The representational framework is warranted only insofar as it identifies consequential distinctions that such methods otherwise risk overlooking.

The proposal admits several potential counterexamples. A correction may be fully propagated while the contested decision remains unchanged for demonstrably legitimate reasons. A claimant may obtain meaningful remedy without any change to the source dataset. An independently competent appeal may operate on primary evidence and never consult the reduced representation. Equally, a richer public dataset may impose disproportionate privacy burdens without strengthening review. Such cases limit any inference that technical transparency, more information, or pervasive update propagation is universally preferable. The ethical adequacy of a specific arrangement remains conditional on its purposes, rights, competent procedures, and actual effects.

The analytical boundary with Paper II is also important. That paper considers contestability and remedy within the broader problem of credible commitments in asymmetric value circulation. The present section focuses on a narrower unit of intervention: the authorization and propagation of corrections to representations employed in consequential decisions. A claim to credibility in a value-return arrangement and a claim to correction of an institutional record can interact, but each requires its own evidential and normative warrant. The boundary with Paper III likewise concerns focus: preservation of normative questions across scales is a property of descriptions under stipulated conditions; contestability concerns the historically situated powers and practices through which descriptions are challenged, revised, and relied upon.

The warrant distinctions formalized in Section 9.10 are especially relevant to the difference between an accepted correction, its documented transmission, and a justified institutional response. They provide a way to report evidential and jurisdictional limits when a correction has no consequential effect.

7.9 Scope of the representational-contestability proposal

The investigation in Section 7 identifies contestable objects, competent reviewers, versioned correction pathways, and downstream reliance as separate dimensions of representational governance. Equation 11 supplies a narrow formal representation of version-sensitive use, while Proposition 7.1 establishes an elementary necessary condition for a correction to propagate through a declared dependency structure. Neither formal result establishes a duty to disclose a given record, reopen a completed decision, or secure a preferred substantive outcome. The table provides an audit vocabulary whose usefulness and limitations require case-specific investigation.

The proposed ethical orientation is to make consequential representational decisions accountable through intelligible reasons, proportionate access to challenge, suitable evidence protection, and authority to address the particular decisions at issue. These aims are open to plural normative interpretation and institutional disagreement. The presence of an appeal form, the implementation of an automated notice, and the publication of version metadata are distinguishable observations; their contribution to justice and effective remedy must be demonstrated with additional evidence. Section 8 examines the accountability conditions associated with consequential institutional use of such representations.

8 Consequential Modeling and Accountability Thresholds

This section examines the responsibilities that accompany the consequential use of institutional representations. Its objective is to distinguish the epistemic warrant for constructing a descriptive model from the institutional warrant for authorizing reliance on that model. The argument proceeds from use-specific accountability profiles and elementary information-preservation limits to authorization, intervention, feedback, divided institutional powers, and revision of the criteria by which claims are recognized. Mathematical results concern specified maps, domains, and rule systems; institutional and normative conclusions require additional warrants. A constructed allocation case provides a comparison of alternative institutional arrangements. The discussion concerns statistical, administrative, qualitative, and computational representations; references to artificial-intelligence governance supply relevant precedents for a subset of these arrangements. The proposed framework is exploratory and establishes neither a general legal standard nor a universal scalar measure of ethical risk.

8.1 Representational inquiry and consequential reliance

A representation may originate as a provisional aid to interpreting records, comparing hypotheses, or revealing regularities. Its use may acquire a different institutional character when a resource allocator, regulator, journal editor, hospital committee, or environmental body relies on the representation to classify a claimant, prioritize intervention, or justify a decision. An analytical report can be consulted without acquiring formal authority, and a formally adopted decision rule can operate with substantial practical discretion. The ethical demands at issue therefore depend on the relationship between the representation, the receiving institution, and the decision pathway rather than on the technical sophistication or stated intent of the model alone.

Four questions are useful for distinguishing these arrangements. First, what claim is the representation warranted to support, and with what limitations? Second, which institutional actor is authorized to act upon that claim? Third, which persons, groups, relations, or ecological conditions can be affected by the action? Fourth, how can the institutional use be questioned, corrected, suspended, or reconsidered? The preservation conditions of Section 4 principally address the first question. The contestability mechanisms of Section 7 concern the fourth. Consequential deployment requires additional attention to authority, causal pathways, and the practical accessibility of safeguards.

The literature offers several adjacent precedents. The NIST Artificial Intelligence Risk Management Framework describes a voluntary, lifecycle-oriented arrangement organized around governing, mapping, measuring, and managing risks [31]. The OECD's AI Recommendation, amended in May 2024, includes a transparency provision encouraging context-appropriate information that enables adversely affected people to challenge an AI system's output; it also addresses human rights, robustness and accountability [36]. This recommendation is a policy instrument, with applicability and legal status distinct from a directly enforceable individual remedy. Raji and colleagues propose an end-to-end organizational audit of AI development and deployment, including documentation generated at successive stages [40]. ISO/IEC 42001:2023 sets out requirements and guidance for organizational AI management systems [21]. These works inform the administrative vocabulary of this section; their specific scope remains artificial-intelligence systems or the organizations using them. The framework proposed here addresses a wider class of representations and should be evaluated independently for those uses.

An institutional authorization is also distinguishable from ethical vindication. A body may possess the competence to publish or employ a summary under its rules while its evidential omissions, affected populations, or remedial pathways remain contested. Conversely, an analytically accurate estimate may be inappropriate for a particular use because it lacks the specificity required by the decision rule, its deployment exposes protected information, or its authorized users cannot respond meaningfully to a challenge. These distinctions provide the basis for the accountability profile below.

Table 10: Consequential-use dossier: dimensions, evidence, and proposed safeguards. The entries are audit questions rather than universal thresholds.

Dimension	Evidence to examine	Possible safeguard
Reach and dependency	Affected uses, receiving institutions, actual reliance and exclusions	Use-specific authorization and dependency register
Severity and reversibility	Plausible pathways of harm, protected interests, feasible correction periods	Escalation, prospective limits, suspension or remedy
Evidential uncertainty	Provenance, coverage, interpretation, validation domain and counterevidence	Independent review and stated uncertainty bounds
Standing and review	Eligible challengers, barriers to access, competence and reasons for disposition	Accessible and protected routes for contestation
Lifecycle change	Versions, shifted use purposes, monitored incidents and corrective actions	Reassessment triggers and withdrawal criteria

8.2 Dimensions of consequence and institutional dependence

The ethical significance of consequential reliance has multiple dimensions. Reach concerns the population and relational structures exposed to a decision. Severity concerns the reasonably supported magnitude of possible loss or impairment. Reversibility concerns the practical prospects of responding to error, including delays and irreversible ecological effects. Dependence concerns the degree to which decision-makers rely on the representation rather than on additional evidence or human deliberation. Authority concerns the competence to adopt, contest, or withdraw the use. Epistemic uncertainty concerns unresolved questions about measurement, historical admission, causal interpretation, and the adequacy of the represented population. These dimensions need not be commensurable; a small number of people may face a grave irreversible consequence, while a widely circulated educational aggregate may have less immediate decisional force.

For a declared use context u , let $C(u)$ denote a profile of the dimensions relevant to that use. The profile is an interpretive dossier with explicitly sourced components, rather than an intrinsic property of a model. A representation may be used in several contexts, producing several profiles. The same summary can be appropriate for a public overview and inadequate for adjudicating individual standing. Moreover, apparent reversibility can change when a classification is replicated across institutions, when a deadline passes, or when material consequences accumulate. Such developments must be reconstructed from records and justified interpretations rather than inferred from a universal clock or a privileged present.

A useful comparison is therefore a *partial assessment*, supported by declared priority relations among the components. Suppose two use contexts have the same lawful competence and evidential basis. One may expose more subjects to a comparable harm and offer weaker feasible review, in which case the justification for additional safeguards is strengthened. Other contexts may be normatively incomparable: wider reach, stronger reversibility, and more uncertain evidence point in different directions. The purpose of an accountability profile is to make those differences inspectable, while preserving disagreements about the weighting of protected rights, distributive interests, care, and ecological effects. The framework does not automatically transform risk descriptors into moral obligations.

Table 10 introduces an audit interface that links each descriptive dimension with a corresponding evidential inquiry and a possible institutional safeguard. The proposed safeguards remain subject to the legal regime, actual institutional competence, confidentiality constraints, and the reasons supporting their proportionality.

8.3 Accountability profiles and a threshold-identification limit

Institutional policies often summarize risk in a band, color, score, or simplified classification. Such summaries can organize attention, although a single indicator may conceal the distinction on which an

appropriate safeguard depends. The following elementary proposition makes this limitation precise for a stipulated accountability rule. It extends the query-preservation logic of Section 4 to the selection of institutional controls, without adding a new mathematical theory of risk.

Let $\mathcal{U}$ be a set of declared uses of representations, let $\mathcal{S}$ be nonempty, and let $\tau : \mathcal{U} \rightarrow K$ a published threshold classification, and $A : \mathcal{U} \rightarrow \mathcal{S}$ a stipulated use-sensitive assignment of safeguard packages from a set $\mathcal{S}$. Here A represents an independently justified policy proposal or requirement supplied from outside the formalism, and τ is a summary of information potentially relevant to its application. A safeguard assignment operating only on the published indicator has the form $a : K \rightarrow \mathcal{S}$. The adequacy condition for exact assignment appears in Equation 12.

$$A = a \circ \tau. \quad (12)$$

Proposition 8.1 (Safeguard assignment and threshold sufficiency). *Under the stated sets and mappings, a function a satisfying Equation 12 exists exactly when $A(u) = A(u')$ for every pair $u, u' \in \mathcal{U}$ with $\tau(u) = \tau(u')$.*

Proof. If a exists and $\tau(u) = \tau(u')$, substitution gives $A(u) = a(\tau(u)) = a(\tau(u')) = A(u')$. Conversely, if A is constant on each fiber of τ , assign $a(k)$ the shared value of A on $\tau^{-1}(k)$ for k in the image of τ ; values outside that image may be chosen arbitrarily when $\mathcal{S}$ is nonempty. Then Equation 12 holds on $\mathcal{U}$. $\square$

The result is elementary factorization. It identifies a *threshold-insufficiency witness*: two contexts given the same public rating while the declared, independently warranted safeguarding rule assigns distinct controls. For instance, an aggregate impact rating may classify two applications identically, although one concerns exploratory comparison and the other determines eligibility for an irreversible relocation. If their review competence, irreversibility, or protected standing warrants different safeguards, that rating alone is insufficient for exact safeguard assignment. The proposition establishes only representational sufficiency relative to A and the admitted domain. Whether A is justified, whether additional distinctions can be recorded lawfully, and whether a feasible control package exists require separate analysis.

An important practical qualification concerns resource constraints. A real institution may choose a stronger common control for both uses, delegate decisions to a protected reviewer, or refrain from the consequential use where it cannot distinguish cases adequately. Such arrangements can be more defensible than demanding publication of sensitive individual records. The proposition also makes explicit why apparent precision in a risk score does not confer normative authority: the authority and warrant of the underlying safeguard rule remain external inputs.

8.4 Composed Representations and the Persistence of Information Loss

Consequential institutions ordinarily rely on a sequence of descriptive operations. Evidence may be admitted through one procedure, classified by a second body, summarized by an analytical unit, and published or passed to a decision-maker under a further rule. Assessing the accuracy of the published statistic provides only one part of the evidential warrant for the final decision. The objective of this subsection is to connect the manuscript's earlier factorization conditions with the institutional dependencies described in Sections 3, 4, and 7. It establishes an elementary limitation on the recovery of information through successive deterministic reductions, then examines circumstances in which supplementary evidence changes the effective information available for review.

Let H denote a declared comparison space of candidate histories, including histories whose exclusion from routine admission is under examination. A record-generation map $E : H \rightarrow X$ assigns each candidate history its observable institutional record. For the stipulated comparison, X may contain a distinguished "not registered" value; doing so permits formal comparison of histories that produce the same absence of an entry. A coding map $C : X \rightarrow T$ assigns institutional classifications or aggregates, and a disclosure map $P : T \rightarrow Y$ defines the output available to a particular audience. The composed

public representation is defined in Equation 13.

$$M = P \circ C \circ E : H \longrightarrow Y. \quad (13)$$

The arrows in Equation 13 specify transformations of descriptions, not temporal precedence, physical dynamics, or moral causation. Each transformation has a declared domain and codomain; if successive offices use incompatible conventions or partially defined interfaces, the composition requires an explicit restriction or reconciliation. A normative evaluation $J : H \rightarrow A$ is fixed only for the purpose of assessing the declared question. Its epistemic and ethical warrant remains a separate matter.

Proposition 8.2 (Persistence of a question-relevant collision under deterministic reduction). *Suppose $E : H \rightarrow X$, $C : X \rightarrow T$, and $P : T \rightarrow Y$ are total deterministic maps. If two histories $h_0, h_1 \in H$ satisfy $E(h_0) = E(h_1)$ and $J(h_0) \neq J(h_1)$, then J cannot be recovered by a function of $M = P \circ C \circ E$ on H . More generally, the same conclusion holds if $C(E(h_0)) = C(E(h_1))$, irrespective of the values of E before classification.*

Proof. Equality of the outputs at the earlier stage implies equality at every subsequent deterministic stage. Thus $M(h_0) = M(h_1)$ while $J(h_0) \neq J(h_1)$. Any purported function $j : Y \rightarrow A$ with $J = j \circ M$ would assign the same answer to both histories, a contradiction. The second claim applies the same reasoning to $C \circ E$ in place of E . $\square$

Proposition 8.2 is a direct consequence of the fiberwise criterion in Equation 2. It contributes no new mathematical theorem about aggregation. Its methodological significance lies in identifying the *earliest established collision* in an institutional description chain. If a claimant's relevant situation is indistinguishable at the entry stage, improving the accuracy of a later public chart cannot independently recover that distinction. If the distinction is recorded at entry but merged at classification, modifying the publication layout alone cannot restore it. This observation does not imply that the originating institution committed an ethical wrong: an omission may reflect lawful competence restrictions, unavailable evidence, deliberate privacy protection, or a genuinely irrelevant distinction. Establishing the nature of the omission requires additional investigation.

A further qualification concerns protected supplementary evidence. Suppose a competent reviewer is authorized to use a second evidential channel $S : H \rightarrow Z$, which may include testimony, direct observations, or records maintained by a separate institution. The reviewer then observes the pair (M, S) , under access rules different from those governing the public output. Exact recovery of the declared question is possible under the formal model precisely when the joint observation has sufficiently fine fibers. This criterion is stated in Equation 14.

$$(M(h) = M(h') \wedge S(h) = S(h')) \implies J(h) = J(h') \quad (h, h' \in H). \quad (14)$$

Equation 14 expresses the elementary factorization condition for a pair of observations. A supplemental channel can distinguish histories collapsed by a public description, but only to the extent that it contains reliable, admissible, and question-relevant evidence. Neither the existence of S nor its informational sufficiency authorizes its collection or disclosure. The question of who may consult protected records, and on what grounds, belongs to institutional and legal assessment. Contextual restrictions may also prevent a reviewer from transferring the information to a consequential decision-maker even where the reviewer can resolve the underlying evidential distinction.

A hypothetical research-infrastructure allocation illustrates the difference between sequential loss and supplementary recovery. Consider two otherwise similar histories involving a contributor whose work was excluded from a formally registered activity category. In one history the contributor had an authorized route to challenge classification; in the other that route was practically inaccessible. If the initial register contains only accepted outputs, both histories can generate the same entry, category and published aggregate. A later model may calculate the accepted-output total perfectly while remaining insufficient for the procedural-access question. A protected review process able to examine independent submissions and the actual review rules could obtain evidence relevant to that question. Such a process

Table 11: A cross-stage audit for identifying information loss and possible institutional recovery.

Stage	Possible loss or ambiguity	Discriminating evidence	Competence question
Evidence admission	Claim or history leaves no distinguishable entry	Submission policies, refusal records, external testimony	Who can review the admission rule?
Classification	Distinct accepted records receive one category	Codebooks, coding reasons, alternative classifications	Who can revise category meanings?
Aggregation	Relevant relations disappear from totals	Source-to-aggregate lineage, disaggregated review	Is refinement permitted for the declared question?
Disclosure	Protected details are withheld from public output	Audience rules, confidentiality basis, reviewer access logs	Who may examine restricted evidence?
Downstream reliance	Corrected versions do not reach consequential use	Model versions, decision files, notification records	Who may reconsider affected decisions?

might nevertheless reach a justified unchanged allocation, might lack jurisdiction over the allocation itself, or might be unable to corroborate the claim. The constructed comparison therefore identifies investigable differences without stipulating a preferred moral outcome.

Table 11 relates the formal stages to possible documentary tests. It distinguishes the location of an informational collision from the competence to address its consequences. A provenance standard can provide a language for derivations and agent participation [46, 47], but conformity to such a vocabulary establishes neither the truth of recorded accounts nor the completeness of the admitted population. Likewise, the NIST AI Risk Management Framework emphasizes intended uses, contextual mapping, monitoring and accountable roles within its voluntary AI-specific scope [31]; those ideas inform the audit analogy without creating new duties for every non-AI institutional record.

The audit in Table 11 should also register cases where the formal problem is absent. A public aggregate may preserve the precise question that an authorized body is competent to answer. A protected reviewer may have adequate access through an independent route even where public recovery is impossible. A valid institutional determination may be unaffected by an admitted classification correction. These possibilities help separate an information-theoretic diagnosis from an empirical finding of inadequate process or an ethical criticism.

The practical proposal is therefore to document the question under examination, the earliest point at which relevant distinctions become unavailable, the authority over that stage, the permissible evidential alternatives, and the decision pathways that might give a correction consequence. Such documentation may sharpen investigation; it cannot settle competing accounts of normative standing or produce reliable histories from inadequate evidence. It is one component of the wider accountability interface, subject to comparison with existing audit, legal and professional procedures.

8.5 Authorization and evidence before consequential use

The passage from inquiry to consequential use can be documented through a proposed *authorization warrant*. Its central object is a use-specific claim: that a defined representation, under a stated version and evidential domain, can be relied on by an identified institution for a stated purpose, subject to review and correction provisions. The warrant should record the source of the institution’s competence, the

persons responsible for authorization, the intended users and affected claimants, the evidence consulted, the uncertainties retained, and the uses excluded from authorization. Its normative requirements arise from the applicable rules and the ethical argument advanced for the relevant setting; the warrant itself does not constitute a legal permission.

The distinction between technical validation and use authorization matters in practical terms. A model may accurately predict an aggregate measure while possessing little validity for a question concerning individual responsibility. A classification may reproduce historical administrative labels while excluding persons whose claims could not enter those labels. A representation may also support several defensible interpretations, each with different implications for public reporting, eligibility decisions, or remedies. Evidence of predictive fit, agreement among expert coders, or consistency of records is therefore informative but insufficient in isolation for ethically consequential authorization. A cautious institutional record should identify which questions the available evidence actually answers.

The warrant can be prepared using existing organizational procedures where these are relevant. NIST’s voluntary AI RMF emphasizes context-sensitive governance across a system lifecycle; Raji and colleagues emphasize the documentary products of internal algorithmic auditing [31, 40]. The present proposal differs in its emphasis on contested admission of histories and the rights or interests made visible through representational design. For institutions outside the scope of AI governance, ordinary archival standards, professional review, administrative law, or participatory arrangements may supply more appropriate foundations. The need to audit a decision does not establish that an AI management standard governs every such decision.

Where available evidence is incomplete, the authorizing institution can state conditions for limited use, establish a further evidential inquiry, or decline authorization. A use restriction should distinguish the reasons supporting restraint: uncertainties in data coverage, uncertainty in causal effects, lack of institutional competence, inability to provide required safeguards, or a substantiated risk of irreversible harm. These reasons may call for different remedies, and a blanket appeal to uncertainty can itself conceal exclusion or shift burdens onto claimants with little access to evidence.

8.6 Monitoring, revision, and institutional withdrawal

Deployment is a continuing institutional practice whose consequences may change with the population admitted to a procedure, the representations used by downstream bodies, and the available resources for contestation. Monitoring should therefore examine the operation of a representation in its declared decision environment as well as the performance of the representation in isolation. Relevant observations may include the use of superseded versions, differing error burdens, newly recognized claimants, changed evidential conditions, challenges that cannot reach a competent reviewer, and corrections that fail to propagate to downstream decisions. Section 7 supplies the dependency framework for distinguishing a revised source from changed practical outcomes.

An accountability plan can specify triggers for review when a material use expands beyond its authorized purposes, significant uncertainty is discovered, protected claimants identify previously unavailable evidence, or the institution detects outcomes inconsistent with an established constraint. The appropriate response depends on competence and severity. It may include renewed validation, a revised representational policy, a restriction on intended use, independent review, notice to affected subjects, reopening a decision where legally and institutionally possible, temporary suspension, or retirement of the representation. These are alternatives to be assessed, rather than a universal ladder that each institution must implement in the same order.

The OECD AI Principles explicitly connect accountability with traceability, systematic risk management, and mechanisms for intervention where AI systems present undue harms [36]. The NIST AI RMF similarly frames governance and monitoring as continuing practices [31]. Their relevance to institutional representations beyond AI is analogical. A technical system’s retirement may be straightforward compared with the revision of a historical government record; a lawful institutional judgment may continue to have effects after the representation that informed it has changed. These differences require explicit

investigation of the authority, the applicable rules, and the practical consequences of withdrawal.

Monitoring also changes the evidential environment. Publication of an indicator may lead institutions to optimize performance against that indicator; affected parties may alter their submissions; records may be revised retrospectively. The audit should record these interactions and avoid treating later records as though they were collected under unchanged representational conditions. Whenever timing matters, the relevant order should be specified as an event, evidential, procedural, or institutional precedence relation. A universal linear timeline is unnecessary for the formal claim that certain decisions depend on particular accepted versions or evidential states.

8.7 Institutional feedback and endogenous representational histories

An institutional representation can alter the practices that it purports to describe when its outputs are used to allocate resources, reward compliance, admit applicants, or justify intervention. Its ethical appraisal consequently requires an account of the responses that a particular use may elicit. This issue is related to, yet distinct from, the loss of a declared distinction under a fixed reduction. The reduction may remain mathematically unchanged while institutions alter registration practices, claimants redirect submissions, or the relevant histories themselves develop under the incentives and constraints created by the decision. The object of assessment then includes both the adequacy of a descriptive rule for a stated question and the effects of authorizing that rule as part of a consequential practice.

Existing scholarship supplies essential precedents for this concern. Campbell warned that the consequential use of quantitative social indicators can create pressures to alter the production of those indicators; his argument was advanced within programme evaluation and should not be read as a universal empirical law [8]. Espeland and Sauder describe *reactivity* in public rankings, including mechanisms of commensuration and self-fulfilling prophecy, in a study of law-school rankings [13]. Their subject matter does not independently establish the same mechanisms in every field. It nevertheless provides a reason to treat the circulation and institutional use of measurements as potential components of the system under study. Work on sociotechnical abstraction similarly cautions against evaluating a technical representation independently of the social context in which decisions are made [43]. The incremental object examined here is the record of how an authorized use changes the conditions under which later claims and measurements become available, together with the limits of the evidence needed to attribute that change.

Conditional continuation and the authority to intervene. Let k and ℓ index two institutional configurations for which $k \prec \ell$ means that decisions in configuration ℓ depend on specified decisions or records in configuration k . The relation $\prec$ is a declared dependency relation; it does not presuppose a universal clock, a total order, or an ontologically privileged present. Let H_k be a contextually admitted set of candidate histories, $R_k : H_k \rightarrow Y_k$ a representation, and $\pi_k : Y_k \times C_k \rightarrow A_k$ an authorized decision rule conditional on recorded institutional context C_k . For a particular relationship between configurations, let $\mathcal{T}_{k\ell}$ associate a history and a decision with a set of admissible continuations in H_ℓ . Equation 15 records the schematic dependence:

$$\mathcal{T}_{k\ell} : H_k \times A_k \longrightarrow \mathcal{P}(H_\ell), \quad h_\ell \in \mathcal{T}_{k\ell}(h_k, \pi_k(R_k(h_k), c_k)). \quad (15)$$

Here $\mathcal{P}(H_\ell)$ denotes the power set, and $c_k \in C_k$ is the stated context for a particular decision. This is a modeling vocabulary, not a theory that decision-makers determine every possible continuation. External events, independent choices, incomplete knowledge, and constraints on institutional authority may enter $\mathcal{T}_{k\ell}$. The correspondence can include more than one continuation and can remain underspecified. The notation becomes useful only when the investigator states what the relevant histories, decision powers, and dependencies mean in the particular setting.

The representational policy at ℓ can itself change. A claimant might learn how a form classifies supporting evidence; a service provider might give priority to cases that increase its reported completion rate; a public authority might introduce a new field following criticism; or an independent reviewer might

Table 12: Proposed discriminating evidence for institutionally reactive representations. The same measured change can admit competing explanations.

Candidate response	Evidence to seek under authorized access	Alternative explanation or limitation
Substantive improvement	Independent outcome traces, practice records, and predeclared outcome definitions	Concurrent resources or policy changes can produce the outcome
Selection of easier cases	Intake denominators, rejected submissions, referral changes, and eligibility rules	Demand and population composition may change independently
Administrative recoding	Versioned forms, coding guidance, reconciliation logs, and source-to-output comparisons	Legitimate corrections may improve the fidelity of the record
Burden displacement	Independently sourced ecological or community evidence across relevant boundaries	Attribution requires exposure, comparison, and causal analysis
Anticipatory behavior	Evidence on knowledge of incentives, decision authority, and changes in routines	Similar change may occur without awareness of the measured target

prohibit a particular use of an index. These acts can alter the domain admitted at ℓ , the effective coding rule, the evidence accessible to reviewers, or the institutional purpose of a previously published output. A comparison of R_k and R_ℓ is therefore underdetermined until the investigator identifies the policy versions, admission rules, population composition, and relevant interpretations. Equal indicator values in two configurations need not mean equal histories, and unequal indicator values do not by themselves establish deterioration or improvement.

Four mechanisms requiring separate examination. The first mechanism is *adaptation within the represented practice*: an institution changes substantive conduct in response to a measure, perhaps by improving care, conserving resources, or reducing waiting time. Such adaptation can be the intended benefit of a well-designed indicator. The second is *selective admission or composition*: cases receiving attention, registration, or service change, potentially leaving harder-to-register claims outside an apparently improved performance rate. The third is *recording and classification response*: a revised form, altered coding convention, or administrative delay changes what becomes measurable, without necessarily changing the underlying condition to the same degree. The fourth is *displacement across institutional or ecological boundaries*: an apparent gain in the evaluated unit coexists with costs transferred to other people, processes, or shared environments. None of these mechanisms is entailed by a numerical trend. The investigator must establish the relevant distinction and seek evidence that discriminates among alternatives.

Table 12 proposes evidential comparisons. The table is a design aid and contains no observation from an actual institution.

A limitation on inference from the deployed policy. An additional epistemic obstacle arises when the decision rule helps determine which outcomes are observed. The observation available under one deployed rule need not identify the consequences of a materially different rule. Let $\mathcal{W}$ denote a declared class of possible institutional response models, including relevant histories, recording arrangements, and sources of uncertainty. For a policy π , let $O_\pi : \mathcal{W} \rightarrow \mathcal{Y}$ denote the observations accessible under its deployment. For an alternative authorized policy ρ , let $V_\rho : \mathcal{W} \rightarrow \mathcal{V}$ denote an evaluative quantity under that counterfactual policy, with its meaning and normative relevance independently specified. The notation is intended to make the assumption of counterfactual comparability visible; it does not assert that the quantity is observable or uniquely determined by the real-world evidence.

Proposition 8.3 (Single-policy observation and counterfactual non-identification). *If there exist $w_1, w_2 \in \mathcal{W}$ satisfying $O_\pi(w_1) = O_\pi(w_2)$ while $V_\rho(w_1) \neq V_\rho(w_2)$, then no function $g : \mathcal{Y} \rightarrow \mathcal{V}$ can satisfy $V_\rho = g \circ O_\pi$ throughout $\mathcal{W}$.*

Proof. If such a g existed, equality of $O_\pi(w_1)$ and $O_\pi(w_2)$ would imply equality of their images under g . By the proposed factorization these images are respectively $V_\rho(w_1)$ and $V_\rho(w_2)$, contrary to the stipulated difference. $\square$

This is the elementary fiberwise factorization criterion already used elsewhere in the manuscript, applied to a different declared query. It is neither an original statistical identification theorem nor a substitute for causal inference. Its contribution to the present exposition is to require the investigator to state what remains unidentifiable after a policy has shaped the observation process. Identifying an effect may call for additional assumptions, independent data, variation in policy exposure, or an ethically permissible comparative design [20]. Such designs must address interference, selection, concurrent reforms, and the possibility that an institution alters its classifications in anticipation of the study. If an alternative policy is legally unavailable or ethically impermissible, the investigator cannot treat unrestricted experimentation as the default method of resolving uncertainty.

Constructed case: a service-completion indicator. Imagine a hypothetical publicly supported service that is evaluated by the proportion of registered applications closed within a stated interval. The indicator is initially used for internal description and later included in a resource-allocation rule. After the change, the published completion rate rises. This observation permits several incompatible interpretations. Additional resources may have enabled genuinely faster and more effective service. Administrative staff may have changed what counts as a completed case. Potential applicants may have been redirected into unrecorded channels. More difficult claims may have been delayed before formal registration. A reduction in time per case may also have affected the quality of assistance. Each interpretation describes a different relation among institutional use, the production of the record, and the conditions experienced by affected people.

Suppose two constructed response models yield the same registered case counts, published completion proportions, and authorized administrative documents under the adopted rule. Under a proposed alternative allocating support by an independently assessed measure of need, the models differ: in one, additional support is effective for high-need cases; in the other, a structural obstacle outside the programme’s remit prevents comparable benefit. The observation set under the adopted policy fails to distinguish these possibilities. Proposition 8.3 states the formal limitation once the two models and policy-specific quantity have been stipulated. It does not adjudicate the competing substantive accounts or establish that an alternative allocation is morally superior. Those assessments require independent evidence about need, effects, rights, authorized purposes, and feasible alternatives.

A responsible audit could compare versioned intake and completion definitions, access protected records through qualified reviewers, obtain voluntary and safe accounts of attempted applications, and study independent service outcomes where they can be evaluated without undue disclosure. It should record negative findings, including a case in which independently corroborated improvements explain the rising rate and there is no evidence of selective exclusion. Investigators should avoid attributing strategic intention to staff or applicants solely because an indicator and a behavior changed together. People may adapt rationally, cooperatively, defensively, or without knowledge of the measured target. Their legal obligations and ethical claims require separate examination from the inference that a response occurred.

Implications for authorization, contestability, and audit design. A use-specific authorization warrant can address reactivity by specifying the intended behavioral effect of a measure, the population and claim pathways it covers, the scope of foreseeable incentives, the external outcomes against which the measure will be checked, and the authority competent to act when its use changes the evidence. Review may call for a revised indicator, protected qualitative evidence, independent outcome monitoring, a changed decision rule, additional administrative resources, or suspension of an inappropriate use. These

possible responses should be appraised for their own burdens. For example, adding a disclosure field may worsen the safety of claimants, while continuous surveillance intended to detect administrative gaming may impose unjustified intrusions on workers or recipients.

The proposed cross-stage audit should distinguish three kinds of change: a change in what occurred or became possible in the relevant practice, a change in the recording and interpretation of that practice, and a change in institutional authority to use the resulting record. These changes may interact and need not share a single temporal sequence across actors. The case-specific dependency record should identify where such interactions can be investigated and where uncertainty must be retained. Accountability therefore concerns the governance of representational feedback, not simply the maintenance of a technically accurate indicator. Appendix B adds prospective safeguards for testing this extension against a mature audit comparator; the added procedure remains unimplemented.

8.8 Distributed responsibility, answerability, and institutional succession

Consequential representations are often assembled and used by people and organizations with different powers. One office admits an account, a second maintains a coding rule, another publishes an indicator, and an independent decision-maker may rely on the publication. A later body can inherit the operative system while having neither the authority nor the source evidence available to its predecessor. The resulting ethical problem is not exhausted by asking who made the representation inaccurate. It also concerns who was empowered to specify its uses, who could reasonably learn of a defect, who was answerable for the decision to rely upon it, and who can initiate an effective response when these capacities are distributed. This subsection develops a *responsibility-interface audit*: a way of posing and documenting those distinct inquiries across an institutional pathway. It proposes a research instrument, not a theory that assigns blame or legal liability through graph structure alone.

The underlying difficulties have important precedents. Thompson's account of the "problem of many hands" examines the difficulty of ascribing moral responsibility to individual public officials when many contribute to a policy, and cautions against treating hierarchical or collective attribution as sufficient substitutes for evaluating individual causal and volitional conditions [45]. Nissenbaum identifies related barriers in computerized systems, including distributed production, treating bugs as inevitable, shifting blame to the computer, and separating ownership from accountability; that analysis concerns computing and does not by itself establish wrongdoing in a particular administrative institution [32]. Bovens conceptualizes public accountability as a relation between an actor and a forum empowered to question, assess, and potentially impose consequences; the existence and quality of such a relationship require institutional examination [6]. Young's social-connection account of structural injustice directs attention to shared, prospective political responsibilities that should not be conflated with backward-looking findings of fault or legal liability [50]. These perspectives illuminate different questions. The present manuscript draws on their distinctions while limiting its own proposal to the accountable circulation of representations and the conditions of their consequential uptake.

Separate registers of responsibility. An investigation should distinguish at least five registers. A *contribution record* identifies activities and transformations attributable to specified participants, with the available degree of evidential confidence. A *role-competence record* identifies the rule, delegation, or agreement under which a body can admit evidence, authorize a model, request reconsideration, or change a consequential decision. An *answerability record* identifies a forum to which an actor must provide reasons and the powers that the forum can actually exercise. A *causal-interpretation record* evaluates whether an activity contributed to a particular effect, including alternative explanations and limits on counterfactual comparison. Finally, a *normative-disposition record* gives independently supported reasons for any criticism, obligation, prospective repair, protection, or allocation of burden. The registers may be linked, but none determines the others by definition.

For example, an archival custodian may be able to explain the provenance of a category without having authorized its use in funding decisions. An allocating body may be answerable for employing a

valid published category even if that body could not edit the source records. A reviewer may have access to protected evidence but lack jurisdiction to reopen a final decision. These are institutionally different facts. Similarly, being a causal contributor to an adverse outcome does not automatically establish a breached duty; possessing a procedural duty to investigate does not establish causal responsibility for the underlying loss. A public-facing organizational chart that describes all participants as “stakeholders” does not resolve these distinctions.

A typed interface rather than a single attribution score. For a declared consequential use u , let S_u be the set of representational or decisional operations whose treatment is material to the question under review. These might include admission, classification, aggregation, purpose authorization, decision reliance, protected review, correction transmission, and implementation of a disposition. The set is inquiry-relative and may be contested. Let I denote the candidate institutions or roles, and F the candidate accountability fora. Define a documentary relation

$$\mathcal{D}_u \subseteq S_u \times I \times F \times E \times K, \quad (s, i, f, e, k) \in \mathcal{D}_u, \quad (16)$$

where $e \in E$ denotes a specified evidential basis (or an expressly recorded limit) and $k \in K$ denotes the stated context of competence, including version, purpose, access rules, and any pertinent jurisdiction. Inclusion in $\mathcal{D}_u$ means only that a putative actor–operation–forum connection has been recorded with its warrant for examination. It does not mean that the actor is culpable, that the forum has authority, or that an evidential claim has been established. A valid *answerability route* additionally requires evidence that i was the relevant actor for s , that f can demand or receive an account under the specified arrangements, and that the route is available for the question and access conditions at issue. The investigator must classify these requirements separately, including *unknown*, *contested*, and *inapplicable* outcomes.

A weak documentary coverage condition for the declared critical operations can be written as

$$\text{Cov}(u; \mathcal{D}_u) = 1 \iff \forall s \in S_u \exists (i, f, e, k) [(s, i, f, e, k) \in \mathcal{D}_u \wedge \text{ValidRoute}_u(s, i, f, e, k)]. \quad (17)$$

Here ValidRoute_u abbreviates explicitly declared documentary and competence checks; it is not calculated from mathematical provenance alone. Equation 17 is a definition of *coverage under a proposed audit*, not a sufficiency criterion for justice or an original mathematical result. Its value can be assessed only for a defensibly delimited S_u . If the set omits a consequential operation because the inquiry itself excluded affected claimants, reported coverage may be misleading. Conversely, a gap in the available documentation may be caused by lawful retention limits or justified confidentiality rather than by the absence of a responsible institution. The underlying reason must be examined before criticism is drawn.

The audit separates two further properties. *Response capability* asks whether a competent forum can order, request, or transmit a warranted response; *consequential reach* asks whether that response can influence the actual dependent decisions within applicable procedural limits. A pathway may provide an explanation and a recorded disposition while having no power to change the downstream use. Another may permit reconsideration but require an independent evidential threshold. The propositions about dependency and correction in Section 7.5 remain relevant: an information path is necessary under their specified assumptions, but it cannot create a missing legal or institutional power. Thus high documentary coverage, genuine answerability, and effective remedy should be evaluated with different indicators.

Table 13 organizes these distinctions as an inquiry into documented roles, institutional powers, and potential response channels.

Institutional succession and the discontinuity of competence. When an institutional configuration k has a documented dependency relation $k \prec \ell$ to another configuration, the transfer of a representation does not entail the transfer of every related power, obligation, or evidential access right. A successor might receive an operational database but be unable to consult protected testimony used by the earlier reviewer. An archive may survive when the programme that generated it has ceased; an appeal forum may continue to exist while its jurisdiction excludes certain concluded decisions. Conversely, a successor

Table 13: Proposed responsibility-interface record. The entries are questions for verification, not findings of culpability or universal legal duties.

Operation	Role and authority evidence	Forum and response evidence	Failure or uncertainty to record
Admission and coding	Submission rules, category version, delegated discretion	Competent review of intake or category choice	Excluded account; protected or unregistered evidence
Use authorization	Intended use, delegating instrument, adopted version	Forum able to demand reasons for use	Descriptive warrant substituted for decisional warrant
Consequential reliance	Decision file, reliance record, counterevidence	Reviewer competent for decision or procedural reconsideration	Reliance disputed; competing causal explanation
Correction transmission	Versioned notification, receipt and dependency links	Body with power to revise dependent use	Corrected source with unchanged operational decision
Institutional succession	Transfer and retention rules, continuing mandates	Successor forum and residual predecessor duties, if any	Custody transferred without access or review competence

can acquire express duties to preserve records and accept challenges even when it did not participate in creating the contested classification. Which situation holds is a question of governing instruments and actual practice. A versioned *succession interface* should therefore record the objects transferred, the basis and limits of access, competent fora, unresolved claims, and any formally established residual responsibilities. The notation $k \prec \ell$ specifies a dependency supported by institutional evidence, without assuming that all institutional histories occupy one globally ordered timeline.

A succession gap should not be resolved by assigning responsibility to a convenient existing custodian. Nor should absence of an identifiable individual cause an investigator to declare an institution irredeemably unaccountable. Where the operative arrangements offer no competent answerability route, possible responses include seeking an authorized referral, a lawful preservation order, a newly constituted review competence, or a reasoned finding that no remedy is currently available. Establishing any of these responses requires an institutional and legal analysis outside the documentary model. Where procedural finality or third-party rights restrict reopening, a transparent explanation, prospective rule change, or independently authorized compensation procedure may be the only feasible candidates; they are not implied entitlements.

Constructed case: a completed-service indicator across three authorities. Extend the hypothetical service indicator in Section 8.7. Agency *A* initially defines a measure counting cases administratively marked “completed”. Agency *B* adopts the published proportion when allocating programme resources. Contractor *C* supplies records and adjusts its workflow in response to the allocation incentive. Following a reorganization, agency *D* takes custody of the database, while an independent review forum *F* can assess a restricted evidential dossier subject to its mandate. A service user’s report alleges that certain difficult cases were marked completed without meaningful resolution. This is an invented arrangement containing no inference about real offices or actual programmes.

Suppose the category was applied consistently under *A*’s codebook, while *B*’s allocation rule required substantive service resolution. A category audit alone could find no coding error. An investigation of *B*’s authorized purpose might identify a mismatch between the recorded measure and the allocation query. A record of *C*’s changed practice could support inquiry into incentive responses, but any claim that the incentive *caused* the change requires comparisons and evidence distinct from a version trace. Agency *D* might be authorized to correct the stored records yet have no competence to revise awards made by *B*. The forum *F* may be able to evaluate a protected complaint without having power to order public disclosure. These differences generate separate and testable questions for the five responsibility

registers.

The same constructed evidence can support several legitimate outcomes. If B 's authorizing documents explicitly confined the indicator to administrative throughput and independent substantive outcomes informed every award, the alleged mismatch might disappear. If C followed a lawful reporting convention but diverted genuinely difficult cases to a different documented pathway, an apparent decline in resolution could reflect changed classification rather than harm. If F possesses a pre-existing and effective cross-agency referral procedure, the new interface could merely duplicate current practice. Alternatively, an unreconciled transfer between A and D could make a justified review practically impossible despite ample technical provenance. Neither actor labels nor mathematically complete data lineage choose among these possibilities.

Protection, burdens, and the risk of surveillance. Making answerability more visible can itself generate harm. Exhaustive participant-level logs may expose claimants, reveal protected testimony, invite reprisals, or shift proof burdens onto persons with fewer resources. A defensible responsibility audit should identify the minimum evidence required for each forum, apply context-specific retention and access controls, record how challenges can be made without unjustified disclosure, and permit independent scrutiny of the audit itself. Public explanations may properly remain less detailed than protected reasons given to a competent reviewer. The distinction between public reconstructability and lawful protected examination in Section 9.2 applies equally to institutional contribution records. Missing public detail cannot automatically be coded as institutional evasion.

Responsibility should also be assessed without assuming that greater individual traceability is always preferable. A system may need to record a responsible office, the basis of its decision, and a confidential route to review without recording every worker or claimant as an enduring identifiable subject. The institutional purpose, risk of retaliation, evidential value, and available less intrusive measures should govern such design choices. These considerations are especially important where the very indicator being audited has reshaped the conditions under which participants feel able to submit information.

Comparative evaluation and limits. The proposed interface has a defensible incremental contribution only if it makes some consequential question more answerable than an adequately resourced existing procedure. A prospective comparison should investigate whether reviewers can identify the relevant actor and forum, distinguish authority to correct a record from power to reconsider reliance, recognize gaps created by institutional succession, and communicate protected reasons without excessive disclosure. It should separately measure the burdens of reconstruction, errors in attribution, participation barriers, and the practical disposition of independently corroborated claims. Increased identification of putative responsibility without additional competent response, or increased logging without safe contestability, is not sufficient evidence of improvement.

The literature on distributed moral responsibility and public accountability supplies the philosophical and institutional groundwork; the proposal here is narrower. It connects the question-relative transformations of this manuscript to a typed record of operational contributions, accountable fora, access limits, institutional succession, and the authority to act upon a correction. It adds no causal estimate or general moral allocation theorem. Appendix B.6 specifies prospective test questions, including negative cases in which a mature existing audit already provides equal or better answerability with fewer burdens. Section 9 retains the distinction between making institutional relations inspectable and justifying the ultimate normative disposition.

8.9 Representation-setting power and the separation of remedial authority

The responsibility interface in Equation 16 identifies institutions able to give an account of particular operations. A further question concerns the distribution of power over the *rules under which an account becomes actionable*. The authority that defines categories may differ from the authority that applies them, the forum that considers a challenge, and the institution able to change a dependent allocation.

These separations may protect independence and lawful division of powers; they may also produce an arrangement in which a claim is heard, recorded, and found material without reaching an institution capable of responding to it. This subsection analyzes these possibilities through a typed account of representational and remedial reach. The analysis is institutional and conditional: documentary routes provide material for inquiry, while powers, duties, and entitlements must be established independently.

Participation, transparency, and consequential influence. Arnstein’s account of participation distinguishes consultation from more substantial powers to determine plans and programmes [4]. Her normative typology arose in a particular setting of United States public programmes; it is a conceptual warning about the allocation of influence, rather than a general rule that consultation is always ineffective. Fung’s framework similarly treats the selection of participants, forms of communication, and linkage to public decisions as separate design dimensions [17]. The analysis of algorithmic transparency by Ananny and Crawford identifies limits to the assumption that observing the operation of a system is sufficient to understand or govern it [1]. These studies give reasons to ask what power an affected person obtains through a review interface, while leaving the legitimacy of any particular distribution of authority open to context-specific argument.

For the present purpose, a *representation-setting power* is the documented competence to determine an admissibility rule, category boundary, evidential priority, aggregation method, purpose of use, or materiality standard. A *remedial power* is competence to take an action on a reviewed claim, such as correcting an operational record, requiring reconsideration, revising an allocation rule, authorizing compensation under an applicable arrangement, or supplying a reasoned disposition within a recognized mandate. A *claimant capacity* concerns the opportunities, costs, and safeguards through which an affected person can submit evidence, question the operative categories, request a review, and seek response. Each category of power is tied to an identifiable institutional object and a source of authorization. A single organization may hold multiple powers; a multi-institution setting may divide them.

The distinction matters where formal invitation conceals selective authority. A review procedure may invite comments on a person’s placement in an existing category while withholding consideration of whether the category is adequate for the decision. A repository may permit a correction of source data while leaving the applicable eligibility rule outside the reviewing body’s jurisdiction. An independent reviewer may identify an unsupported use of a score yet have authority only to recommend further investigation. None of these arrangements proves that the category, rule, or decision is unjust. Each requires a description of the claimant’s actual options and of the institutional powers relevant to the declared question. The present audit therefore records *scope-setting competence* as a distinct object: who may specify which representations, claims, decision rules, and remedial actions are admissible to review.

A two-layer inquiry into information and authority. Consider a stipulated use u and a claim episode c . Let V_u contain documentary or institutional locations relevant to c , including a submitting route, a custodian, qualified reviewing fora, and possible deciding bodies. Define two directed, context-labelled relations over these locations. The first, L_u^{ev} , records an independently supported route through which a specified evidential object can be transmitted and reviewed subject to lawful access and admissibility conditions. The second, L_u^{auth} , records an independently supported institutional capacity to request, authorize, transmit for implementation, or implement a specified form of response. Edges may overlap, but membership in one relation has no implication for membership in the other. A reviewer’s access to records does not confer appropriation authority; a funding authority’s budget power does not establish competence to examine protected testimony.

For a forum f , deciding body d , and proposed action type a , introduce three predicates under declared contextual and evidential assumptions. $\text{EvReach}_u(c, f)$ means that the claim can lawfully reach an appropriately competent evidential assessment. $\text{ActReach}_u(f, d; a)$ means that there is an authorized institutional route from that assessment to possible consideration of a by d . $\text{Adm}_u(c, f, d; a)$ records the independently justified admissibility, standing, access, and procedural conditions for that particular

combination. A narrow *potential remedy reach* criterion is then defined as follows.

$$\text{PotReach}_u(c, d; a) := \exists f \in V_u [\text{EvReach}_u(c, f) \wedge \text{ActReach}_u(f, d; a) \wedge \text{Adm}_u(c, f, d; a)]. \quad (18)$$

Equation 18 is a definitional audit test for a particular proposed response, not an entitlement or a guarantee of implementation. An authorized channel to *consider* compensation differs from a power to award it; and an institution authorized to reopen a decision may have sound evidential, legal, or ethical reasons for maintaining it. Actual remedial action, its substantive effect, its effect on third parties, and its justification must each be recorded independently. A finding of $\text{PotReach}_u = 0$ may identify a documented route gap, an indeterminate access condition, or an inapplicable requested action. Inquiry into these alternatives should precede any normative conclusion.

Proposition 8.4 (Documentary indistinguishability of remedial competence). *Fix a claim c , a use u , and a proposed response a . Suppose two admissible institutional configurations share the same observable evidence-transfer relations, source versions, and recorded claim dispositions, while one configuration contains a duly authorized route from the reviewing forum to a body competent to consider a and the other does not. Then any function of the shared documentary observations alone has the same value in both configurations, although their potential remedy-reach predicates in Equation 18 differ. Consequently, the existence of a remedial route is not identifiable from those observations without further information about institutional powers and applicable conditions.*

Proof. Let the common documentary observation be o . A function $h(o)$ has the same value in the two stipulated configurations. By hypothesis the authorized route and the other conditions required by Equation 18 hold in one configuration and fail in the other. Thus the value of $\text{PotReach}_u(c, d; a)$ differs while o is held fixed. No function of o alone agrees with both values. The assertion is a direct indistinguishability argument under the stated hypotheses, and supplies neither a new graph theorem nor a factual determination about any institution. $\square$

Proposition 8.4 provides a warning about audit evidence. Two organizations may publish indistinguishable correction logs, complaint counts, provenance diagrams, and procedural descriptions even though only one has a duly delegated route to revise a downstream decision. Conversely, the existence of a formal route may be documented while the substantive review remains ineffective because of resource constraints, inaccessible evidence, insufficient notice, or other barriers. A complete study should independently establish the scope and version of a mandate, any budgetary and execution constraints, the relevant review window, and whether a protected evidential record can be used without exposing persons who reasonably require confidentiality. Following the manuscript's treatment of relational histories, institutional dependencies may be partly ordered by documented transfers and authorization relations; a common chronological sequence is not assumed by this definition.

Critical separations and contestability of the scope itself. Where an evidential route reaches a reviewing forum but no authenticated authorized route exists for the relevant response, the audit may record a *remedial disconnection*. A stronger, graph-specific term, *authority cut*, may be used for a declared response and declared finite network if removing a stated set of authorized connections disconnects every admissible route to the deciding body. These descriptions require the investigator to identify the network boundary and the response type. They cannot imply that every action outside the chosen network is impossible, or that an institution is under a duty to create the missing connection. New competence may require legislation, contractual amendment, appropriations, agreement between bodies, or an entirely different process of political deliberation.

A particularly consequential separation arises when the same body that specifies admissible categories also sets the initial limits of their contestation. Let Q_u denote the declared family of questions to which a decision rule is purportedly responsive and let Γ_u be its formally permitted set of review questions. A challenge about the adequacy of a category can be material to Q_u while absent from Γ_u . This generates a *scope challenge*: a dispute over whether the authorized review itself recognizes the question

Table 14: Separated representational and remedial powers in a proposed case audit. Each row specifies distinct evidence and a different limit on inference.

Inquiry	Evidence to obtain	Question for a competent forum	Unresolved limit
Representation setting	Codebook, exclusion rules, approval record, purpose of use	Who may revise the category or its permitted uses?	Evidence of version control does not establish fairness
Claim reception	Notice, standing, intake and protected referral rules	Can a challenge reach qualified assessment safely?	Registered complaints omit claims never submitted
Scope challenge	Review mandate and admitted question family	Can the adequacy of the rule itself be examined?	A classification appeal may exclude rule design
Decision reconsideration	Decision instrument, review powers, third-party constraints	Who may reconsider the dependent decision?	Reconsideration may properly confirm the decision
Implementation capacity	Execution and funding instruments, time limits, documented follow-up	Can an authorized disposition have consequential effect?	Formal authority differs from capacity and outcomes

that matters. Recording a scope challenge requires no assertion that all questions deserve the same forum, or that public disclosure of sensitive records is appropriate. It requires an express account of who can revise Γ_u , how excluded questions can be received safely, and what competent external or democratic avenues are available, if any. Where no such avenue is identified, the limitation should remain visible rather than converted into an apparently completed appeal.

The cost of exercising a formally available option also matters. Scholarship on administrative burden identifies learning, compliance, and psychological costs that can shape access to public programmes [19]. An appeal pathway may have legal standing and an empowered reviewer while remaining practically inaccessible to a claimant who cannot learn of it, afford its procedures, or risk disclosure without substantial harm. A proposed *effective-access review* would therefore investigate accessible notice, translation where relevant, reasonable documentation requirements, safeguards against retaliation, representation or assistance where authorized, and opportunities for anonymous or protected preliminary inquiry. These are context-specific assessment factors, not universal legal entitlements inferred from the cited literature.

Table 14 separates five investigations that can otherwise be compressed into a single “access to appeal” indicator. It also provides an explicit negative case: if a standard institutional review already verifies all five questions and reaches the appropriate decision-makers without additional burden, the augmented interface may provide little incremental value.

Constructed case: an inter-agency needs allocation. Consider a hypothetical regional programme that distributes support among local service providers. Statistical office R maintains a versioned needs classification; allocating board B uses an aggregate index to assign funds; local office L receives individual accounts and can correct its source entries; independent reviewer F can assess protected complaints and issue recommendations. The rules authorize R to modify the classification prospectively, B to determine allocations under the adopted scheme, and F to review individual intake errors. No particular national statute, actual programme, or claimant is represented. A group of affected persons raises a hypothetical concern that a type of hardship omitted from the index has become relevant to the programme’s stated objective. The individual source records are accurate according to their current definitions, but the measure may be incomplete for that objective.

Suppose F can receive a confidential account and verify that L entered each permitted field accu-

rately. It then closes the complaint as a correctly recorded case. This outcome can be procedurally proper under F 's limited mandate while leaving the challenge to the index's domain unresolved. The audit would document the narrow scope of the disposition, rather than count it as an unsuccessful factual correction or as vindication of the index. If F may refer a material category challenge to R , and R can lawfully revise the index for future use, one potential response route exists. Yet R 's power to revise a codebook may not empower B to alter an earlier allocation, and the procedural status of completed awards might exclude retrospective revision. Independent legal and institutional evidence determines the permissible form of response.

Several counter-cases discipline the inference. The omitted hardship may be addressed effectively by another authorized programme, rendering its omission from this index consistent with the stated allocation purpose. A broader index may impose disclosure risks that justify a protected qualitative supplement rather than public addition of a field. An institution may have made a reasoned, participatory decision to constrain the programme's scope under legitimate resource limits. Alternatively, affected persons may have no practical way to challenge the category and no competent forum authorized to consider it. The records alone do not rank these configurations ethically; a defensible conclusion requires evidence about purposes, participation, costs, authority, outcomes, third-party claims, and available alternatives.

An expanded audit could therefore preserve a versioned *scope-challenge dossier* containing the declared question, the affected representational distinction, the route by which evidence was received, the independently verified mandates of F , R , and B , the possibilities and limits of protected access, the stated grounds of the disposition, and any lawfully authorized follow-up. It should not publish personal narratives by default or infer misconduct merely from institutional fragmentation. The dossier's practical test is whether it helps distinguish a legitimately limited procedure from a material representation dispute that remains outside all effective routes, while maintaining appropriate protections.

Institutional design, duties of care, and conditional remedy. Possible design responses include an authorized referral between category-setting and decision-review bodies, explicit standing to seek review of a declared materiality standard, independent examination of protected evidence, publication of the permitted scope of a review, or a prospective review point for changed uses. Each proposal has costs and institutional implications. A more powerful single reviewer might weaken specialized expertise or procedural independence, while fragmented competencies can protect privacy and distribute checks. A revision that reallocates resources can affect people whose original decisions were lawful; procedural finality may serve legitimate reliance interests. Normative assessment should therefore address proportionality, legal competence, justice, the distribution of burdens, and available alternatives, with clear distinctions among prospective reform, factual correction, reconsideration, compensation, and explanation.

The proposed interface is strongest when it makes a *missing institutional question* visible and identifies a lawful way to ask it. Its weakness is that the diagnostic vocabulary alone cannot create power, guarantee voice, determine which claims should prevail, or repair harms. The combination of a verified authority map, a protected evidence route, and a declared category-scope inquiry is accordingly a subject for comparative evaluation rather than a complete theory of institutional justice. Appendix B proposes separate measurements of scope challenge, potential reach, practical accessibility, and substantive outcome. The larger framework treats the distribution of representational control as part of the ethical problem while retaining independent standards for deciding which institutional transformations are justified.

8.10 Continuity of contestability across institutional transformations

This subsection examines the continuity of a materially specified challenge when records, representational rules, or institutional competences are transferred or revised. Its objective is to identify the conditions under which a claimant's question remains intelligible and institutionally addressable across several authorized operations. The proposed model treats a handoff as a context-dependent relation among

claims, records, and powers. It does not assume a universal sequence of institutional states, uninterrupted retention of personal information, or a general right to reopen an otherwise final decision. Its first concern is the distinction between a technically reconstructible record and a practically available procedure.

Representational lineage and competence lineage. Let K be a set of documented institutional configurations. A configuration identifies a representational version, declared purposes, record custody, evidence-access restrictions, review fora, and powers over specified responses. A directed link $k \rightsquigarrow \ell$ records an independently corroborated relation of transfer, derivation, delegation, or institutional succession relevant to a particular inquiry. The relation denotes a recorded dependency or authorized change of configuration; it does not posit that the labels in K exhaust institutional history or define an ontologically privileged chronology. Different branches may coexist, different authorities may retain overlapping jurisdiction, and a relation established for one issue need not apply to another. The historical and legal basis of each link belongs in the evidential record.

Two lineages should be reconstructed separately. *Representational lineage* concerns the derivation of a category, score, source document, aggregate, or redacted record. *Competence lineage* concerns the mandate to admit a challenge, consult protected material, revise a classification, reconsider a dependent decision, or implement a specified response. A version-control entry can document the former without establishing the latter. Conversely, an authorized successor reviewer may possess competence to hear a new submission without acquiring the predecessor’s restricted records. The provenance vocabulary in W3C PROV-DM can describe entities, derivations, agents and bundles of provenance [46], yet a provenance graph supplies no independent legal warrant for evidence disclosure or for remedial action. The accountability and authority distinctions introduced in Sections 8.8 and 8.9 remain operative at each configuration.

Typed claim states and guarded handoffs. For each $k \in K$, let C_k denote the *analytic* set of claim states that an inquiry considers at that configuration. A state specifies a question about a represented history, a permitted evidential locator, its status of corroboration, the available protected-access route, the relevant version of a rule, and a requested response type. It may use a pseudonymous reference or a qualified independent intermediary; it need not contain, or publicly reveal, a claimant’s identity. These are categories for an audit, not presumed data that an institution is licensed to collect or retain.

For a documented link $e : k \rightsquigarrow \ell$, define a relation $R_e \subseteq C_k \times C_\ell$ containing precisely those pairs for which a claim can be carried, reconstituted, or appropriately referred under the established conditions governing that link. Membership calls for an inspectable warrant concerning scope, evidence access, applicable time limits, authorization and communication of the route to the claimant or qualified representative. An unknown warrant should be recorded as *unresolved*, rather than silently converted into a positive or negative edge. The relation may be empty even if a database transfer has succeeded. It may also be populated through a protected independent submission despite the absence of any transfer of identifiable records.

For a finite, contextually ordered handoff path $p = (e_1, \dots, e_n)$, composition is defined only where the successive claim-state types match. The path relation is stated in Equation 19.

$$R_p = R_{e_n} \circ \dots \circ R_{e_1} \subseteq C_{k_0} \times C_{k_n}. \quad (19)$$

A *witnessed continuation* for an initial claim state $c_0 \in C_{k_0}$ consists of a path, intermediate claim states, and the independently corroborated edge warrants needed to establish $(c_0, c_n) \in R_p$. A *potentially actionable continuation* additionally requires a terminal forum qualified for the material question and a separately verified potential-response route for the proposed action, as in Equation 18. These properties are indexed by claim and requested response. They establish availability of a route under stated assumptions; they neither predict a favorable outcome nor determine whether that outcome would be ethically justified.

Table 15: Proposed evidential fields for auditing a challenge across institutional transformations.

Audit object	Evidence or distinction to record	Possible continuity failure
Question and version	Material issue, category definition, intended use, version relation	Scope changes without a route to question the applicable rule
Protected locator	Lawful identifier, referral token, independent-submission route, access and deletion conditions	Source and receiving claim states cannot be matched safely
Handoff authority	Transfer, delegation, review competence, standing, limitation and finality rules	Valid record transfer is mistaken for transfer of jurisdiction
Evidential adequacy	Independent corroboration, restrictions, uncertainties, reviewer qualifications	Available metadata cannot establish a material contention
Possible response	Eligible forum, permitted action, third-party interests and implementation route	Review conclusion cannot reach a competent response body
Actual disposition	Timely notice, costs, response, substantive reason, consequences and disagreement	Formal referral masks abandonment, burden or adverse effect

Proposition 8.5 (Local handoff availability and global claim continuity). *For a path of two or more transfers, nonempty local handoff relations do not suffice for a witnessed continuation of a particular initial claim. A witnessed continuation exists for that claim if and only if there is a compatible sequence of intermediate claim states belonging to each consecutive relation. A qualified terminal forum and potential-response route are additional conditions for potentially actionable continuation.*

Proof. Consider $C_0 = \{a\}$, $C_1 = \{b, c\}$ and $C_2 = \{d\}$, with $R_1 = \{(a, b)\}$ and $R_2 = \{(c, d)\}$. Both handoff relations are nonempty, but $R_2 \circ R_1$ is empty because no intermediate state participates in both relations. More generally, the definition of relational composition in Equation 19 makes a pair (c_0, c_n) a member of R_p exactly when a compatible sequence of intermediate states exists. The additional conditions for a qualified terminal forum and response route concern institutional competence, which is not contained in the composition alone. The result is an elementary relational observation, without a claim of mathematical novelty. $\square$

Proposition 8.5 identifies a failure mode obscured by stagewise compliance summaries. A sending institution may refer a challenge under a source identifier, and a receiving institution may operate a valid route for a new classification, while the receiving route accepts only a different identifier or only objections to data-entry mistakes. Counting both routes as “available” will not establish the viability of the initial category-scope challenge. The mismatch might be corrected through lawful linkage, interpretation, or referral. Alternatively, the competent response may lie in an independent forum or the original decision may properly be final. These are distinct documentary and normative findings.

An evidential dossier for continuity. Table 15 proposes fields that permit comparison of such findings without conflating a recorded routing edge with a claim’s actual outcome. Its categories are deliberately capable of recording protected access, justified closure, and indeterminate authority. The documentation should be proportional to the seriousness of the disputed use and the expected burdens of review.

The dossier in Table 15 should also record who may contest a categorization of the challenge itself. A receiving authority might accurately handle all accepted correction requests yet define a broad objection to the admissibility rule as outside its process. In that case the issue is the distribution of scope-setting power already examined in Section 8.9. The audit may reveal a missing procedural bridge, but identifying a missing bridge cannot itself confer competence upon another institution.

Constructed case: a reclassified service programme. Consider a stipulated service programme administered through three organizational configurations. In configuration *A*, an agency uses a category labelled “completed service” for internal measurement and accepts an independent submission challenging which forms of support qualify. In configuration *B*, the operational register is migrated to a shared

provider, and a revised codebook narrows the category to digitally confirmed transactions. The migration preserves each aggregate count and a provenance entry. In configuration *C*, a separate allocating body uses the revised count to determine eligibility for a supplementary programme. This is a constructed example, without an empirical allegation about any organization.

Under one possible arrangement, the provider's correction route accepts only mistaken transaction identifiers; the agency's earlier category-scope submission is held under a protected reference that the provider cannot lawfully inspect; and the allocating body can review arithmetic consistency but lacks authority to reconsider the category rule. All three organizations may operate valid procedures within their own mandates. The original challenge nonetheless has no documented compatible claim-state route to a forum capable of considering the requested response. The loss is partly representational, partly jurisdictional. Neither the absence of a direct data transfer nor a claimant's failure to repeat a protected account establishes personal negligence or institutional wrongdoing.

Under a second arrangement, a qualified independent reviewer can receive the earlier concern, verify the pertinent version and category rule through permitted evidence, and refer a reasoned finding to a body independently authorized to determine whether reconsideration is permitted. The claim may proceed without disclosure of protected testimony to the provider or public. This arrangement creates an investigable possibility of response, yet a justified decision to retain the allocation remains possible. Under a third arrangement, a binding finality rule or applicable privacy duty may lawfully preclude reopening or retaining a particular link. Where the forum provides an adequately reasoned disposition within those limits, the audit should register lawful closure, alternative prospective correction if available, and any independently contested normative questions. It should not classify closure automatically as injustice.

Confidentiality, erasure, and legitimate noncontinuity. A permanent, universally linkable claimant identifier would often increase surveillance and disclosure risk. The proposed continuity test therefore permits protected tokens, limited-purpose referral proofs, access by a qualified reviewer, prospective category revision without identity transfer, or a new independent claim with adequate evidence. These arrangements should be tested for feasibility rather than assumed universally available. Under the EU General Data Protection Regulation, for example, processing is subject to purpose limitation, data minimization, storage limitation and specified rights to rectification and erasure, each with legally relevant conditions and exceptions [14]. That legal framework supplies a jurisdiction-specific illustration of potential constraints; it does not govern every institutional setting described here. Contextual privacy obligations and relational expectations also require attention beyond technical identifiers [33].

A justified break in a particular claim's documentary continuity may protect a vulnerable person or third-party interest. Evidence-preservation and retention policies should therefore be justified relative to a legitimate review purpose and available alternatives, with minimal access and demonstrable deletion or compartmentalization safeguards. A demand for more complete lineage can itself worsen safety, compromise source confidentiality, or suppress participation. The question is which distinctions a qualified process needs to examine, who may access them, and what burden a person must bear to make the process usable. The administrative-burden literature motivates separate attention to learning, compliance and psychological costs [19]; their prevalence and magnitude in a given programme require empirical investigation.

Comparative value and limits of the proposed continuity interface. The continuity relation is a structural audit device. Its empirical merit depends on whether reconstructing compatible handoffs detects independently material concerns that existing professional, archival, legal or administrative procedures fail to address. The comparison should measure correct referrals, competent review of scope challenges, safe preservation of admissible evidence, errors in assumed jurisdiction, claimant burden, actual dispositions, and consequences for affected third parties. It should count negative findings, including cases in which the incumbent process is already adequate, a purported gap concerns only an inapplicable response, protected nonretention is warranted, or added linkage produces excessive exposure and no

improvement. The prospective module in Appendix B.8 specifies an evaluation design. This inquiry extends the documented representational and competence lineages without deciding how far legal powers or duties should extend across institutions.

8.11 Ethical priority among competing challenges and remedies

This subsection investigates the ethical evaluation of competing representational challenges once their evidential pathways and potential remedial authorities have been identified. The problem differs from the continuity problem in Section 8.10. A fully documented, competently received claim can give rise to several potentially authorized actions, each affecting different persons, institutional capacities, and protected interests. An institution may also receive multiple independently credible claims within the same decision setting. The practical question is how to make a defensible selection, ordering, or deferral without assuming that information recovery itself supplies an ethical ranking. Three judgments must be distinguished: *material seriousness* concerns the significance of a representation for a specified question; *review urgency* concerns the reasons for giving a challenge procedural attention within a declared horizon; and *remedial justification* concerns which available response is permitted and warranted. These judgments can diverge.

The present analysis draws a limited connection to the literature on procedural priority setting. Daniels and Sabin’s account of accountability for reasonableness specifies public justification, relevance, revision and appeals, and regulation or enforcement as conditions associated with the legitimacy of resource-allocation processes [10]. Its original setting is health-care priority setting. The conditions are relevant to the design of institutional procedures considered here, but they neither establish a universal legal entitlement nor decide which claimant or institutional response deserves priority in every domain. Work on plural values and incommensurability questions the assumption that heterogeneous goods can always be represented on a single scale [44, 2]. The present manuscript’s independent, narrower contribution is to make the limits of such comparisons inspectable within its question-indexed documentary interface.

Separation of procedural urgency and remedial merit. Suppose a documented challenge concerns a misclassified applicant. Its severity depends in part on the consequential use of the category, whether the error has been independently corroborated, and the burden plausibly associated with the classification. Its review urgency may depend on an independently established deadline, a credible risk of irreversibility, or the safety of awaiting a decision. Its requested remedy may entail individual correction, reconsideration by a competent decision maker, prospective revision of the category, compensation, or restricted disclosure. Neither urgency nor materiality establishes that any one response is legally available, proportionate, effective, or just. A procedure might properly expedite a request for protected investigation while deferring any determination on consequential relief. Conversely, a documented high-impact structural issue might justify broad consultation and prospective revision, even where no claimant-specific urgent action is appropriate.

Nor should a queue be mistaken for a theory of moral worth. A category such as “urgent” is a procedural assessment indexed to a decision context, access conditions, and possible delay-related effects; it is neither a declaration that another person has fewer rights nor a discount factor on historically significant injuries. Institutional triage should publish or securely disclose to qualified reviewers the reasons and evidence for its procedural distinctions, specify reconsideration conditions, and monitor patterns of systematic non-entry. A procedure that prioritizes only well-documented cases may systematically disadvantage persons for whom recording a claim is costly or unsafe. Conversely, an unconditional demand for more detailed records can expose protected testimony or impose further administrative burdens [19, 11].

A constrained and uncertainty-aware response set. Let $q \in \mathcal{Q}$ denote a materially specified, independently eligible question, which may concern one or several claimants. Let A_q be a declared set of

candidate institutional responses. Inclusion in A_q means only that an action is under consideration; it does not confer legal power or a favorable evaluation. For the purpose of a prospective audit, let Ω_q be a documented set of evidentially admissible interpretations or decision scenarios that remain open after qualified review. The set is a methodological device for representing uncertainty, not an assertion of exhaustive states of reality, fixed probabilities, or an ontologically privileged present. For an action $a \in A_q$ and scenario $\omega \in \Omega_q$, an independently verified institutional feasibility and rights-compatibility screen records a three-valued status. Equation 20 specifies the possible statuses.

$$F_q(a, \omega) \in \{\text{permitted, excluded, unresolved}\}. \quad (20)$$

The screen incorporates the relevant mandate, protective restrictions, established standing and finality requirements, and applicable safeguards for others' rights. Its entries require legal and institutional evidence. It has no force to resolve those matters by definition. The *verified feasible set* is defined in Equation 21.

$$A_q^{\text{vf}} = \{a \in A_q : F_q(a, \omega) = \text{permitted for every } \omega \in \Omega_q\}. \quad (21)$$

This conservative screening convention is suitable for a particular kind of robust documentary comparison; it may be too strict for institutions authorized to make decisions under specified evidential thresholds. Such alternative thresholds must be justified independently and disclosed. An *unresolved* feasibility assessment is not evidence of prohibition, and an *excluded* action cannot acquire authority through favorable projected outcomes.

For feasible actions, the audit may record a vector $v_q(a, \omega) = (v_{q1}, \dots, v_{qm})$ of contextual, adverse-effect descriptors. Potential dimensions include avoidable residual injury to directly affected persons, collateral interference with third parties, disclosure or retaliation risk, reversibility concerns, delay-related burdens, and opportunity costs of scarce institutional capacity. Each dimension must have an explicitly stated referent, interpretation, direction of comparison, evidential basis, and uncertainty. Some entries may be ordinal and others may be qualitative. Not all can be reduced to a shared quantitative unit. A vector is therefore a documentation device for locally comparable components, rather than a utility representation of justice or a vector of objectively measured harms. An investigator must not invent numerical values for experiences that have not been independently assessed.

Where, and only where, the selected dimensions are defensibly comparable, write $a \prec_q b$ if $a, b \in A_q^{\text{vf}}$ and every justified adverse-effect comparison weakly favors a over b across the admissible scenarios, with at least one strictly favoring a for a documented dimension and scenario. In a fully numerical illustrative case this becomes

$$\begin{aligned} a \prec_q b \quad &\iff \quad \forall \omega \in \Omega_q, \forall j \in \{1, \dots, m\} : \\ &v_{qj}(a, \omega) \leq v_{qj}(b, \omega), \\ &\text{with strict inequality for some } (\omega, j). \end{aligned} \quad (22)$$

Equation 22 is a conditional componentwise dominance test with conservative uncertainty treatment. The ordinary mathematical ordering is inherited, not proposed as a novel theorem. When the components or the admissibility of cross-person comparisons are disputed, the audit should record *incomparable* or *unresolved* rather than manufacture a ranking. The strict relation does not apply to prohibited actions, and a comparison can change when the justified dimensions, admissible evidence, or institutional mandate change. A robustly non-dominated set may provide a defensible starting point for deliberation, but its members are not thereby morally equal, legally authorized beyond the screen, or collectively optimal.

Proposition 8.6 (Incomplete priorities do not establish a unique ethical response). *Assume two verified feasible actions are assessed only through the componentwise test in Equation 22, with two equally admitted adverse-effect dimensions and one admissible scenario. The test need not select a unique action. A unique total ranking requires an additional principle or further justified comparison not supplied by feasibility and componentwise dominance alone.*

Proof. Let $A_q^{vf} = \{a, b\}$, $\Omega_q = \{\omega\}$, $v_q(a, \omega) = (1, 0)$, and $v_q(b, \omega) = (0, 1)$, where smaller values are locally preferable on each dimension. Neither vector weakly improves upon the other componentwise. Thus neither $a \prec_q b$ nor $b \prec_q a$ holds. An ordering of these actions would require an additional trade-off, priority, or decision procedure. The example establishes a limitation of the stated test, not that informed ethical deliberation can never reach a justified decision. $\square$

The limitation in Proposition 8.6 matters because the demand for a single published “remedial priority score” might reproduce the representational compression being criticized. For example, a weighting rule assigning equal cardinal coefficients to disclosure risk and economic injury would already make a substantive judgment about the commensurability of distinct goods. A regulator may have valid reasons to adopt such a rule for a narrowly described statutory purpose. The ethical task is to articulate its authorization, safeguards, justificatory reasoning, and sensitivity to reasonable disagreement. The manuscript does not prohibit aggregate indicators; it rejects the assumption that an aggregate alone establishes justification. Where legal rights operate as constraints, they should be identified before optional cost comparisons rather than priced as routine penalties that an advantageous score can overcome.

Capacity, delay, and the rights of absent parties. Let $\mathcal{B}_k$ describe an institution’s verified capacity constraints within an explicitly stated operational decision horizon at configuration k . A review plan π may assign personnel, protected evidential access and decision slots to a subset of claims. Its *procedural feasibility* requires confirmation that the necessary review powers, expertise, protective conditions and resources exist. This is a different question from whether a specific remedy is warranted. The effect of choosing one plan on other eligible claimants, including those not yet visible to the institution’s intake system, belongs in the distributional assessment. Capacity scarcity may warrant transparent scheduling, institutional escalation or provision of assistance; it does not by itself eliminate duties or prove the legitimacy of denying a claim.

A documented urgency criterion may use the credibly supported risk of irreversibility under an explicitly specified continuation, subject to evidence and appeal. The relevant horizon is an institutional modeling choice, not an assertion of universal temporal ordering. Admissible alternative continuations may differ in their evidential status. Their omission should be declared where it changes the priority result. An urgent temporary protection can be distinguished from a final finding of responsibility, and authorization for expedited review from authorization for disclosure, compensation, or reopening.

Third-party considerations require particular care. Correcting an aggregate record may reveal information about persons who did not bring the challenge; reallocating an authorized resource may affect others who relied upon the prior rule; reopening a concluded file may implicate protected finality interests. These effects should be examined through proportionate, authorized review. A right of one party to submit evidence is no general permission to access another party’s records, and a claim of systemic importance does not suspend confidentiality obligations. Conversely, confidentiality must not be used as an unexplained code for dismissing a claim when protected independent evaluation is possible. The auditor should record which restriction is supported, whether a lawful substitute route exists, and which otherwise plausible responses are ruled out.

Constructed case: competing corrections in a public-service programme. Consider a hypothetical programme whose published “completed case” indicator is used to inform resource allocation. A person P claims that their unresolved request is misrepresented as completed and seeks urgent individual correction because a documented programme deadline may affect access to a service. A second person Q requests broader publication of the source records, arguing that case-level evidence is necessary to contest the indicator. A third group R challenges the underlying definition of completion, contending that a workflow can conform to the codebook without providing substantive resolution. The programme custodian can amend an input record, an independent reviewer has conditional access to protected testimony, and a different authority controls resource-allocation criteria. All parties and facts are constructed.

Table 16: Illustrative distinctions for evaluating competing challenges in the constructed service programme.

Object	Review question and evidence	Finding to keep separate
Case-specific correction	Source codebook, protected case evidence, consequential dependency	Material error versus entitlement to alter a dependent decision
Review urgency	Documented decision horizon, reversibility, safe access	Expedite inquiry versus prejudge its disposition
Requested publicity	Data-access powers, privacy interests of nonclaimants, substitute protected review	Warranted scrutiny versus authorized release of personal information
Category-scope dispute	Rule-setting mandate, purpose of the indicator, affected groups' reasons	Correct coding versus appropriateness of the category
Resource competition	Verified staffing, alternative routes, eligible but unregistered claims	Capacity limits versus moral ranking of persons
Competing responses	Rights screens, scenario uncertainty, adverse-effect dimensions	Robust dominance, incomparability, or unresolved warrant

The three questions require different kinds of inquiry. The credibility and materiality of P 's claim may justify expedited protected verification; it does not automatically justify altering an allocation decision made by another authority. Q 's requested publication may expose the personal information of nonparticipating service users; a secure audit by an independent authorized reviewer is a distinct candidate response, contingent on a valid legal basis. R 's category-scope challenge may warrant prospective consideration by the competent rule-setting body, even if no data-entry error is present. A pending review of R 's challenge need not prevent a permissible temporary protection for P . The effects of prospective reallocation on nonparticipating beneficiaries must be examined separately from the correction of the underlying indicator. These actions might be compatible, in tension, or institutionally independent. A single numerical queue cannot establish their proper relation.

Table 16 illustrates why a credible record may support an action to investigate while remaining insufficient for one or more requested remedies. It also clarifies that the ability to inspect protected evidence may enlarge the options for deliberation without producing a general entitlement to wider publication. The case is not evidence that a particular institutional configuration works, nor a benchmark from which a quantitative improvement estimate can be inferred. If an existing review process already distinguishes these pathways and protects third-party interests at reasonable cost, the proposed interface could add administrative complexity without relevant benefit.

Procedural justification, revision and restrained closure. A defensible process should record a question-specific rationale for the admitted options, the powers of the choosing institution, the permitted access to evidence, pertinent reasons that were offered and considered, the role of absent or less visible parties, and the routes to qualified reconsideration. Reasons should be communicated at an appropriate level of disclosure, including protected channels when warranted. Daniels and Sabin's procedural conditions offer an important comparative precedent, while their health-care setting and the present manuscript's separate rights and representational questions limit direct transfer [10]. The evidence of serious competing obligations may justify limited interim measures, targeted additional inquiry, a reasoned decision among incomparable options, or a bounded refusal. Recording incomparability is compatible with making an accountable institutional decision; it prevents the audit model from claiming to have derived that decision from an incomplete metric.

Comparison with established practice should assess whether the proposed distinction between urgency, justification and remedial competence changes what a qualified forum can understand or protect. It should also test whether the extension delays urgent safeguards, increases burdens, conceals decision authority behind technical notation, or privileges easily documented injuries. The prospective module in Appendix B.9 treats adverse or equivalent results as substantive findings. The normative conclusion of this subsection is deliberately limited: ethically consequential reductions require a record of their interpretive and procedural assumptions, while actual priority decisions require independent authority,

substantive reasons and contestable judgment beyond the representations used to support them.

8.12 Governance of evaluative frames and recognition of affected parties

The preceding subsection assumes that a question q , a feasible response set and the pertinent dimensions of evaluation can be declared for scrutiny. Those declarations are themselves consequential representational acts. A framework that carefully assesses competing claims after admission may continue to exclude persons whose experiences fall outside its intake categories, or may fail to register a challenge to the criteria by which importance, urgency and admissibility are defined. The present subsection therefore investigates *frame governance*: the institutional processes through which relevant questions, affected parties, evidential categories, evaluative dimensions, and review powers are adopted and revised. Its object is neither a universally complete list of affected persons nor a single criterion for all legitimate decisions. It is the inspectability and contestability of consequential choices about the boundary of the evaluative inquiry.

This concern has several intellectual precedents that should be kept distinct. Fraser’s discussion of abnormal justice highlights contestation over the frame within which claims of justice are made, including who counts in such a frame [15]. Anderson’s account of the epistemology of democracy addresses the role of diverse participation, dissent, and feedback in collective learning [3]. Ostrom’s discussion of polycentric action identifies possibilities for learning and initiative across differently situated decision centers [37]. These works differ in object, normative foundation and domain; none establishes that a particular committee, invitation rule or versioned registry is democratically legitimate. They motivate a bounded institutional question: when a consequential representation is contested, what makes revision of its evaluative terms epistemically credible, procedurally accessible and answerable to an authorized forum? The account of actor–forum accountability in Bovens [6] and the differentiated participation model in Fung [17] provide related, but nonidentical, points of comparison.

A typed evaluative frame. For a declared institutional configuration k , let

$$\Gamma_k = (\mathcal{Q}_k, \mathcal{S}_k, \mathcal{D}_k, \mathcal{V}_k, \mathcal{A}_k, \mathcal{J}_k), \quad (23)$$

where $\mathcal{Q}_k$ records questions that the institution is prepared to consider; $\mathcal{S}_k$ records the operative standing and submission rules; $\mathcal{D}_k$ records the permitted evaluative dimensions and associated evidential warrants; $\mathcal{V}_k$ records the procedures for reviewing those dimensions and decisions; $\mathcal{A}_k$ records candidate actions and their governing authorization screens; and $\mathcal{J}_k$ records the institutions and forums competent to make or review decisions. Equation 23 supplies the audit vocabulary. Each component is a *typed documentary specification*, not a set of morally correct answers. Legal competence, admissible evidence, and justified norms require independent assessment. The index k labels an institutional configuration in a declared comparison; it does not presuppose a natural, universal succession of periods. Where there are parallel or contested versions of a rule, their comparative relation may be partial.

For a specified material question q , the frame supplies an admission relation $\text{Adm}_{\Gamma_k}(c, q)$ for a possible claim or claimant c . The relation may be *admitted*, *excluded*, or *unresolved*; this three-way recording is necessary because missing forms, unwillingness to disclose an identity, and an independently established exclusion are different evidential conditions. An accepted-case record $L_{\Gamma_k}(h)$ is produced from a declared candidate history h . Such a record can support claims about processing *within* the institutional frame, but it does not exhaust the candidate histories or persons that independent inquiry might identify. Let C denote a stipulated domain of possible claimants for the declared question. We accordingly distinguish the institution’s recorded claimant set $C_{\Gamma_k}(h) \subseteq C$ from an independently ascertained, question-indexed set $E_q(h) \subseteq C$ of potentially affected persons or interests. The latter is an evidence-dependent study construct. It may remain incomplete or partly protected and need not be made into an identifiable public database. Nor does potential affectedness, by itself, establish the right to a specific disclosure, decision or remedy.

A formal gap indicator may be written

$$G_{q,\Gamma_k}(h) = E_q(h) \setminus C_{\Gamma_k}(h). \quad (24)$$

Equation 24 describes a *conditional ascertainment gap* relative to a declared independent assessment, rather than a measure of the entire public or a direct finding of injustice. Its interpretation depends on the definition of affectedness, lawful ascertainment, evidential sufficiency, and the possibility that some persons appropriately decline to enter the register. A protected person may be counted in the independent analysis without disclosing their identity to the deciding institution. An institution may also lack competence to hear some matters that nevertheless merit acknowledgement or authorized referral. These cases should be coded separately from wrongful exclusion.

Proposition 8.7 (Accepted-case records do not identify completeness of recognition). *Fix a declared question q and frame Γ . Let H be a stipulated comparison space of candidate histories, $L_\Gamma : H \rightarrow Y$ the resulting accepted-case records, and $G_{q,\Gamma} : H \rightarrow \mathcal{P}(C)$ the independently defined ascertainment-gap sets. If there exist $h_0, h_1 \in H$ such that*

$$L_\Gamma(h_0) = L_\Gamma(h_1), \quad G_{q,\Gamma}(h_0) \neq G_{q,\Gamma}(h_1), \quad (25)$$

then no function of accepted-case records alone determines the gap set on H . This is a conditional non-identification result, not an assertion that the actual gap has been established.

Proof. Suppose that $f : Y \rightarrow \mathcal{P}(C)$ satisfied $G_{q,\Gamma} = f \circ L_\Gamma$. Equality of $L_\Gamma(h_0)$ and $L_\Gamma(h_1)$ would imply $G_{q,\Gamma}(h_0) = G_{q,\Gamma}(h_1)$, contrary to Equation 25. Hence no such f exists. $\square$

The result is an elementary instance of the same fiberwise dependency test developed in Section 4; no mathematical novelty is claimed. It also illustrates why a high acceptance rate or successful disposition rate among already admitted claimants cannot demonstrate comprehensive access. An off-register sampling strategy or authorized independent review may supply additional evidence, but its own frame must be stated, including who could not safely be contacted and who may be ineligible for reasons justified independently of the challenged category. The proposition neither entails a universal obligation to register everyone nor determines the substantive merits of an excluded person's claim.

Frame-changing operations and review competence. A transition $\Gamma_k \rightsquigarrow \Gamma_\ell$ may involve at least four distinguishable operations. A *record correction* changes an evidential entry while maintaining the operative category; a *scope change* modifies which questions, persons or accounts can be admitted; a *criterion change* changes admissible evaluative dimensions, thresholds, or their interpretation; and an *authority change* reallocates control over rulemaking, verification or remedy. The same institutional episode may combine operations. The comparison must record which type occurred, which versions were used to decide pending cases, the reasons given, and the competent body involved. Merely versioning a document does not make a change legitimate. Likewise, correcting an erroneous data entry does not in itself answer a challenge to the original definition of the variable.

Let $\rho_{k\ell}$ denote an explicitly documented, possibly partial correspondence between elements of Γ_k and Γ_ℓ . It may map a category to a narrower set of categories, map a question to an authorized successor forum, or record that no lawful identity linkage may be retained. There is no presumption that $\rho_{k\ell}$ is invertible, that every claim survives the transition, or that earlier decisions must be reopened. The correspondence is useful for asking whether a qualified reviewer can reconstruct which operative frame generated a consequential outcome and can identify the authorized rules for applying a changed frame. An absent correspondence can indicate lost traceability, a valid protective separation, or incomplete knowledge. Evaluation must determine which interpretation is warranted.

The possibility of *self-certifying scope* requires special consideration. Suppose the only persons able to challenge a rule defining affectedness are those whom that rule already recognizes. The procedure may be internally coherent and offer meaningful reconsideration for admitted claimants, yet lack any

Table 17: Prospective audit of the authority and evidential basis of evaluative-frame revision.

Revision object	Documentary and independent inquiry	Limit that must remain visible
Recognition boundary	Admission rule, off-register ascertainment, protected participation possibilities	Unregistered is not synonymous with wrongfully excluded
Declared question	Material dependency of the category on consequential decisions	Another question may require another legitimate forum
Admissible dimensions	Origins, epistemic limits, dissent and effect of proposed revisions	No metric establishes its own moral authority
Revision authority	Legal or institutional mandate, reasons, review and notice channels	Documentation cannot create competence
Version transition	Rule text, effective scope, migration rule, case-specific protection	Later categories need not retroactively govern earlier files
Third-party safeguards	Access restrictions, safe proxies, indirect harms and refusal rights	Recognition does not confer unrestricted disclosure rights
Outcome and countercase	Actual disposition, competing explanations, existing independent review	Added procedure may yield no benefit or cause harm

mechanism to hear a reasoned challenge to its own admission rule. That is a vulnerability in the structure of contestability, not proof that each exclusion is unjust. A qualified independent scope reviewer, an authorized representative with a protected route, a public-law review mechanism where applicable, or a procedurally justified external consultation can supply different possible solutions. Their availability, legitimacy and actual accessibility must be documented rather than presumed. Requiring every frame to be open to every external challenge would be equally indefensible, particularly where privacy, relevance, jurisdiction or safety properly limit participation.

Constructed case: a revision of a service completion indicator. Extend the hypothetical programme in Section 8.11. Its rule Γ_0 defines a completed service request as one for which a provider has entered a closure code. Registered applicants can correct mistaken codes; the panel authorized to reconsider the allocation formula accepts testimony from registered applicants. A community advocate submits, through an independently lawful route, an account of persons who attempted to apply but could not complete an intake step. Under one configuration, those accounts are not in the service register and the panel construes its own jurisdiction as limited to discrepancies within the registered population. Under a second configuration, a separate and competent reviewer can evaluate the evidential basis for changing the scope of the completion definition, without demanding public identification of those affected. Neither configuration has been observed; the example is deliberately constructed.

A revision $\Gamma_0 \rightsquigarrow \Gamma_1$ could preserve the original closure-code accuracy while introducing a separately defined *application barrier* indicator, with its own source and uncertainty information. Another reasonable institutional response might refrain from publishing any new count until safe independent ascertainment and consultation are possible. A third could determine that the allocation body lacks the required authority and refer the matter to a competent policymaking forum. The underlying decision to change funding, grant compensation, or preserve existing allocations remains separate. In particular, a newly adopted category may expose absent parties to new harms if it invites inference about identities, turns qualitative experiences into stigmatizing counts, or diverts capacity from urgent assistance. Additional representational visibility is therefore a candidate means of recognition, not a sufficient ethical end.

Legitimacy, epistemic improvement and the possibility of justified closure. Participatory access, epistemic improvement, lawful authority and distributive justice are connected but distinguishable standards. Hearing additional testimony may reveal that an indicator suppresses a relevant distinction; it cannot alone determine the lawful scope of an agency's mandate. A lawful revision procedure may remain epistemically weak when its evidence omits people unable to use official language or register safely. A well-informed and accessible process can nonetheless yield rival defensible normative judgments. The

audit should record whose reasons could enter, whether competent decision makers engaged with them, what disagreement remained, and why an option was adopted, rejected or referred. It should not infer legitimacy from the number of submitted comments, nor invalidity from the persistence of dissent.

Table 17 separates the items to be checked from the authority and harm claims they cannot settle. The study of polycentric coordination offers a further caution: the existence of several decision centers can enable comparison and experimentation, but it can also fragment authority and place referral burdens on persons least able to navigate them [37, 19]. An independent scope channel therefore requires evaluation of actual competence and practical access, together with procedures for reasoned refusal and warranted finality. Epistemic diversity can improve access to information but is neither sufficient for every morally correct judgment nor a substitute for protected investigation [3, 11]. Rights protections and the authority of applicable institutions remain external constraints on the proposal.

This extension establishes a research requirement rather than a definitive institutional design. Where a frame is revised, the record should expose the previous and adopted questions, the admissible evidence and excluded accounts, the relevant institutional powers, the treatment of pending cases, the reasons for departure, and safeguards against displacement or publicity harm. Where a revision is refused, it should distinguish absence of jurisdiction, uncorroborated factual premises, protected reasons, and a substantive decision to preserve the criterion. Comparative evaluation must ask whether these distinctions improve defensible recognition or review beyond mature existing procedures, and whether they add risks, costs or selective visibility. Appendix B.10 specifies such a prospective inquiry. A well-governed evaluation frame remains contestable in principle under appropriate procedures, while particular questions may be legitimately concluded without perpetual reopening.

8.13 Reflexive revision and the limits of self-authorizing evaluative frames

The frame-governance analysis in Section 8.12 requires a further distinction. An institution might permit challenges to its categories and nevertheless subject *those very challenges* to the contested categories. A request to expand standing, change the admissible evidential language, or challenge the mandate of a reviewer may then be rejected for failing a rule whose validity is the object of the request. The analytical problem is not that every instance of self-review is illegitimate. Administrative and professional bodies can acquire lawful competence to reconsider their own policies. The problem is to distinguish a reasoned, authorized exercise of such competence from an internally successful procedure that treats its own eligibility test as conclusive evidence of the test's legitimacy.

The literature offers several different orientations rather than one institutional solution. Wide reflective equilibrium permits reconsideration of judgments, principles and background assumptions in light of one another, although coherence does not by itself establish that those excluded from deliberation have been appropriately considered [25]. Pettit's account of contestatory democracy gives particular importance to inclusive and responsive opportunities to challenge governmental discretion; its constitutional argument is not a ready-made test for the legitimacy of every service-programme committee [38]. Experimentalist governance examines iterative goal-setting and revision among decentralized authorities under specified conditions, without implying that iteration automatically supplies lawful jurisdiction [41]. Lindahl's work on the boundaries of legal order foregrounds the challenge that claims to normative relevance can exceed an order's own categories; the present manuscript uses that observation to motivate questions about institutional boundaries, not to derive rules of positive law [28]. These distinct approaches help frame a constrained inquiry: what additional evidential and institutional relationships can make revision of a contested admission rule examinable without imagining a procedurally neutral standpoint outside every normative frame?

Object-level correction and meta-level contestation. Let Γ be the typed evaluative frame in Equation 23. Introduce an institutional *revision configuration*, defined in Equation 26.

$$\Xi_k = (\Gamma_k, \mathcal{M}_k, \mathcal{U}_k, \mathcal{F}_k, \mathcal{T}_k), \quad (26)$$

where $\mathcal{M}_k$ identifies the types of proposed changes admissible for consideration; $\mathcal{U}_k$ describes procedures for submitting and evaluating challenges to that admission rule itself; $\mathcal{F}_k$ identifies competent forums and the limits of their powers; and $\mathcal{T}_k$ records authorized changes of configuration and their evidential lineage. The index k denotes a documented configuration, not a universal temporal instant. An internal case correction, a revision of Γ_k , an alteration of $\mathcal{M}_k$, and reassignment of $\mathcal{F}_k$ concern different objects. Recording an invitation to submit comments at the first level does not establish access to the others.

Let $\mathcal{C}^*$ be a stipulated universe of possible challenges, including challenges to standing and to the revision mechanism. An institutional state $s \in \mathcal{S}$ includes a declared configuration $\Xi(s)$ and the admissible documentary record. Write $a_s(c) = 1$ when the internal rules admit challenge $c \in \mathcal{C}^*$ to a competent consideration route at s ; a zero can represent a justified rejection, a procedural exclusion or an unresolved access failure, distinctions that must be separately recorded in any application. A directed edge $s \rightarrow_c s'$ denotes a *permitted internal revision attributable to the consideration of c* . Such an edge requires $a_s(c) = 1$ as a bookkeeping condition; admission by itself does not imply that a revision occurs. Other edges may reflect staff turnover, legislation, technical correction or independently authorized review. We shall not silently attribute those changes to an excluded claimant.

For a starting state s_0 , define $\text{Reach}_{\text{int}}(s_0)$ as the states reachable through the declared *internal-only* revision relation. In this comparison the internal relation is generated by challenges admitted under the then-operative rule and excludes changes introduced by an additional lawful channel. This is a bounded model of a particular revision architecture, not an ontology of institutional time or an assertion that institutions follow deterministic evolution.

Proposition 8.8 (Closed-gate non-reachability under internal-only revision). *Fix a challenge $c^* \in \mathcal{C}^*$ and an initial state s_0 . Suppose that*

1. $a_s(c^*) = 0$ for every $s \in \text{Reach}_{\text{int}}(s_0)$; and
2. *any transition explicitly attributable to institutional consideration of c^* must be an admitted-challenge edge $s \rightarrow_{c^*} s'$.*

Then no internal-only path from s_0 contains a transition attributable to the consideration of c^ . In particular, an internal-only revision log cannot establish that this challenge received substantive consideration merely by documenting a sequence of otherwise authorized changes.*

Proof. Every source state along an internal-only path from s_0 belongs to $\text{Reach}_{\text{int}}(s_0)$. By the first assumption $a_s(c^*) = 0$ at each such state. The edge condition in the second assumption therefore excludes an edge labeled by c^* at every step. The claim follows for any finite path; for an infinite path it follows for each finite prefix. Other changes may incidentally satisfy the demand stated by c^* , but they cannot, on these records alone, be attributed to institutional consideration of that excluded challenge. $\square$

This is an elementary graph-reachability result, not a new theorem of constitutional theory. It establishes a boundary of inference under the stated edge-label convention. The premises can fail. A legislature may lawfully revise the admission rule without first admitting the particular challenge; an institution may possess an independent statutory duty to reconsider eligibility; a court or ombud forum may hear scope objections under a different standing rule; or a public consultation may provide an authorized alternative. Nor does the result imply that every excluded challenge is substantively meritorious. It identifies the evidential weakness of claiming *reflexive openness from the internally admitted revision record alone*.

Authorization and justification across recursive levels. A proposed revision $r : \Xi_k \rightsquigarrow \Xi_\ell$ should be accompanied by a differentiated warrant profile, defined in Equation 27.

$$\mathfrak{W}(r) = (W_{\text{evid}}, W_{\text{comp}}, W_{\text{access}}, W_{\text{reasons}}, W_{\text{effects}}), \quad (27)$$

recording respectively the credibility and limits of evidence; the institutional competence to entertain and enact the proposal; safe access by relevant challengers; the reasons for an adopted, rejected or deferred

revision; and observed or reasonably anticipated consequences for affected and third parties. These entries have different justificatory roles. A complete rationale is not the same as valid competence; lawful competence does not settle ethical justification; and an apparently favorable outcome does not retrospectively authorize prohibited disclosure. The profile is a typed checklist of questions, not an additive score, an exhaustive moral standard or a certificate of legitimacy.

Recursive contestability need not entail an infinite series of mandatory meta-appeals. Different legal orders may supply constitutional rules, representative rulemaking, ordinary administrative review, sunset and reconsideration provisions, professional oversight, or judicial remedies. These sources can impose relatively stable but revisable limits on an institution's discretion. Their own legitimacy may be politically or ethically contested without every routine decision requiring a fresh foundational settlement. The inquiry therefore distinguishes *operative procedural validity* within a stated legal configuration, *reasoned normative defensibility* in light of broader affected interests, and *availability of challenge* to a competent forum. A revisable system can possess warranted points of closure, provided that the basis and scope of closure are identified rather than inferred from the absence of recorded objections.

There is no universally external, institution-free observer in this account. Independent sampling, a protected reviewer, and an appeal body each operate under their own rules of evidence, jurisdiction and confidentiality. An additional reviewer may reproduce exclusions, substitute one constituency's priorities for another's, or increase the surveillance of persons who intended not to register. The meaningful question is whether a specific *lawfully competent* route can receive a contested ground and engage with its merits under protections appropriate to the context. A nominal appeal route that can transmit documents but cannot consider the admission criterion does not supply that capacity. Conversely, a competent, accessible existing review mechanism may make an additional multi-level register unnecessary.

Constructed case: an eligibility rule and its revision rule. Consider a hypothetical service office whose completion indicator counts successful outcomes among registered applicants. Its advisory committee permits objections by current registered applicants to entries in the completion register; the same eligibility rule governs proposals to change the register's definition. An association representing persons whose prior service encounters could not be entered under the existing category seeks reconsideration of the eligibility definition. The fictional complaint provides corroborating records through a confidential representative. The internal portal returns an inadmissibility notice because the complainants are not registered. Subsequently, the office updates its reporting format through routine internal review, but the advisory record contains no deliberation on the association's scope objection. The situation satisfies the proposition if the internal-only paths do not introduce an authorized route for that objection. It does not establish that the association is entitled to inclusion, that the old indicator violates law, or that the format change was caused by the complaint.

Three alternative continuations should be compared rather than assumed equivalent. In the first, a pre-existing administrative review tribunal has authority to examine the eligibility rule and provides a protected route; the relevant remedy may already exist without adopting this manuscript's supplemental register. In the second, the programme has power only over clerical records, while alteration of eligibility rests with a separate elected or statutory authority; a valid referral must distinguish information transfer from that authority's decision competence. In the third, confidential evidence cannot lawfully be shared in identifying form, but an authorized independent assessor can verify its materiality under protective conditions and provide a reasoned, appropriately redacted finding. The third route may permit consideration while preserving a person's legitimate decision to withhold public testimony. None supplies an automatic duty to change the rule.

Table 18 records the distinct evidential questions arising from this constructed challenge.

Consequential implications and limits of reflexive governance. An institution may need to revise its rule for determining admissible questions while lacking valid grounds for retrospectively reversing decisions made under earlier rules. Pending claims, third-party reliance, procedural finality, accessible reasons and jurisdictional transitions should therefore be considered separately. Transparent change logs

Table 18: Documentary questions for a constructed challenge to a revision rule.

Inquiry	Required documentary distinction	What the record cannot settle alone
Admission	Object-level claim versus objection to standing or meta-review	Whether the excluded challenge is legally or morally correct
Competence	Referral, rule interpretation, amendment and remedy powers by forum	Whether the forum ought to acquire additional powers
Protected evidence	Public summary, restricted evidence and independent corroboration	Permission for unrestricted access or publication
Revision lineage	Origin of proposal, operative versions, decisions and alternative causes	Causal attribution from chronology alone
Outcome	Hearing, reasoned disposition, rule amendment, consequential effect	Ethical legitimacy from a completed workflow
Closure	Finality grounds, reopening rules, safeguards and alternative routes	Universal entitlement to repeated appeals

may help identify conflicts between successive criteria; they cannot resolve conflicts about rightful authority or distributive justice. Similarly, broader participation is not a sufficient indicator of improved epistemic quality: the impact of participation on reasons, decisions, protective access and materially affected people should be independently examined.

The relevant empirical comparison is with already effective forms of independent review, policy consultation and lawfully authorized amendment. The question is whether a structured meta-challenge record improves the capacity to distinguish internal correction from rule-level reconsideration, improves access to competent dispositions, or reveals harmful closure that other practices miss. Adverse outcomes, including longer delays, increased strategic complaints, greater exposure of protected persons or displacement of urgent casework, would count against the incremental proposal. Appendix B.11 develops a prospective evaluation module with negative cases. Nothing in the present formalism establishes a final, self-validating ethics of evaluative revision; rather, it identifies conditional informational, procedural and institutional questions that a defensible revision practice should be prepared to answer.

8.14 Constructed comparison: a coastal adaptation allocation

Consider a hypothetical regional programme allocating funds for coastal adaptation. The programme produces a public index of exposure based on registered infrastructure assets, recorded hazard reports, and standardized geographical units. Its authors initially develop the index for regional description. An allocating authority subsequently proposes to use the same index to rank local applications. The example is stipulated and bears no claim about any actual authority, community, dataset, or allocation.

In arrangement *A*, the authority treats the published index as a sufficient ranking basis. Local applicants can correct spelling or geocoding errors through the data custodian, but the allocation body has no process for considering the omitted care infrastructure, informal access routes, or locally recorded ecological dependencies relevant to certain applications. A correction to the index may also arrive after an award has been made. The published figures may therefore be arithmetically reproducible even when the suitability of the index for the allocation decision remains unresolved. The weakness lies in the mismatch between its original descriptive warrant and its consequential institutional use.

In arrangement *B*, the allocation body publishes the index's authorized uses and material limits, identifies supplementary evidence admissible for allocation, and provides a competent reviewer for contested classifications. When a correction is accepted, an associated procedure records the versions supplied to decision-makers and determines whether the allocation remains open for consequential reconsideration. Protected evidence may be available to the reviewer while remaining outside the public index. This arrangement is not guaranteed to be more just. Its justification depends on the burdens created, the credibility of the alternative evidence, the reviewer's independence, the rights of other applicants, and applicable funding rules. Its comparative advantage is the availability of identifiable institutional questions for which the two arrangements can be assessed.

An empirical inquiry could compare comparable allocation episodes, reconstruct the admission rules and dependency records, and investigate which claims obtained substantive consideration. It could also examine whether additional review improved decisions or merely created formal process, whether the more elaborate arrangement disproportionately burdened smaller applicants, and whether public administrative resources accounted for observed differences. Claim registration, correction, substantive reconsideration, changed allocation, and improved ecological resilience would need separate measures. The hypothetical example therefore motivates discriminating research rather than demonstrating causal effectiveness.

The claim-specific adjudication matrix in Section 9.10 separates representation adequacy, institutional reliance, legal competence, and normative evaluation. In applications of the coastal allocation example, these distinctions constrain conclusions drawn from any observed record correction.

8.15 Competing explanations and research limits

A proposal for heightened representational accountability should be tested against simpler institutional explanations. Resource scarcity, substantive statutory restrictions, decision-maker expertise, existing appeal rules, organizational incentives, and the quality of independent monitoring may each explain institutional outcomes. An apparently beneficial audit may work because it increases staffing rather than because it improves representational preservation. Conversely, a formal review procedure may exist yet be ineffective where claimants lack access to documents or where authorized revisions cannot affect subsequent decisions. A comparative study must distinguish these possibilities before crediting a representational intervention with improvements in justice, trust, or ecological conditions.

The proposed thresholds also face genuine normative disagreement. Some accounts prioritize non-interference with valid institutional decisions; others emphasize procedural inclusion, historical repair, care, precaution, or the standing of affected ecological interests. The present framework can make the consequences of those choices more inspectable and can reveal when different authorizing interpretations require different evidence. It cannot settle their relative moral authority by applying Equation 12. Protected information imposes further limits: the information needed by a competent reviewer may be inappropriate for unrestricted publication. Appropriately separated access levels and reasoned decisions may be better justified than either total disclosure or a wholly opaque process.

The framework's explanatory value is also conditional. Where a representation is used only for openly non-consequential illustration, additional procedural infrastructure may impose costs without improving relevant epistemic or normative outcomes. Where a pre-existing legal or professional process already provides appropriate evidential review, a new representational audit may duplicate protections. Where the underlying moral disagreement concerns substantive entitlements rather than representational adequacy, improvements in a model's fidelity may leave the core dispute unresolved. These are possible limits on the contribution claimed by this paper and appropriate negative cases for future investigation.

The principal conclusion is a bounded methodological one. Consequential reliance calls for a defensible account of the representation's authorized use, its evidential adequacy, the decisions and subjects it can affect, and the institutional response available when its limitations become material. Proposition 8.1 explains why a summary indicator may be insufficient for safeguard assignment relative to a declared warranted rule. Table 10 specifies an investigable documentation interface. Neither supplies a self-executing legal obligation or an objective scalar ethical threshold. Section 9 examines hard cases in which well-documented institutional representations remain contested and identifies limits of the research programme.

9 Hard Cases, Limits, and Research Agenda

The final section evaluates the limits of the proposed ethics of representational reduction and specifies a research agenda capable of testing its practical contribution. The preceding analysis has examined scale selection, admission of histories, normative information preservation, historical annotations, epistemic

exclusion, contestability, and consequential modeling. The objective here is to place those components under pressure in cases where their requirements conflict or where the information needed for their application is incomplete. The discussion proceeds through conflicts among defensible interpretations, protected information, institutional competence, historical and ecological uncertainty, and competing accounts of model failure. It then audits the scope of the mathematical results, separates the proposed incremental contribution from established scholarship, and specifies comparative tests against mature existing procedures. All cases in this section are hypothetical analytical constructions.

9.1 Interpretive disagreement and representational sufficiency

A representation can be sufficient for a declared question while remaining controversial in the institutional practice that makes the question consequential. One research team may regard allocation accuracy as the relevant evaluative object; another may identify participation in the allocation procedure as a protected interest; a third may emphasize historically transmitted ecological responsibilities. Their questions can require different evidential variables and different interpretations of the same record. A useful representation for one inquiry may preserve too little for another, even where its numerical calculations are sound.

The difficulty concerns the authority to define the family of questions. In Sections 4 and 3, adequacy depends on stipulated histories, interpretation conventions, and queries. In a consequential institution, the choice of these stipulations is itself a matter of governance. A committee could announce a narrow family of admissible questions and certify a compact summary as complete relative to that family. Mathematical sufficiency would then establish a fact about the declared evaluation task. The committee's authority to exclude a competing question, and the ethical defensibility of doing so, would require separate argument. The distinction matters even if the committee follows its own procedures correctly.

Consider a hypothetical repository that summarizes community contributions to a research infrastructure through validated transactions and accepted outputs. A distributive interpretation assesses whether the recorded benefits are proportional to those outputs. A procedural interpretation asks whether contributors may inspect a rejected classification and seek consequential review. An ecological interpretation asks whether infrastructure operations place enduring burdens on surrounding communities. A transaction summary may support the first query but provide little information about the latter two. The appropriate conclusion is task-relative adequacy for the distributive query, coupled with documented uncertainty about broader normative adequacy. The example does not establish that the alternative interpretations possess equal moral or legal authority. The institutional inquiry must examine their grounds, standing, applicable rules, and supporting evidence.

A possible response is to retain an *interpretation register* identifying the evaluations for which a reduced description was designed, the principal disputed alternatives, the evidence each alternative would require, and a reasoned procedure for changing the register. Such a register would expose the selection of questions to criticism. It would not make disagreement disappear, and it could itself become an object of strategic capture. Institutional settlement may be necessary for action under disagreement while leaving reasons for dissent and claims to further review in circulation. This proposal draws upon scholarship concerning classification and the construction of institutional categories [7], while its specific review procedure remains an exploratory proposal of the present paper.

9.2 Confidentiality, publicity, and review access

The preservation of relevant information and the legitimacy of its circulation impose distinct requirements. Consider two possible institutional records that agree on public aggregate measures, but differ in the identities and circumstances of claimants whose submissions warrant confidential treatment. A public summary may protect these claimants by omitting sensitive attributes. The same omission may prevent an outside observer from independently reconstructing eligibility decisions. An institution that seeks both secure disclosure practices and effective contestability must therefore specify the channels by

which a competent reviewer can inspect additional evidence, subject to applicable privacy obligations and procedural safeguards.

Nissenbaum’s account of contextual integrity directs attention to the norms governing information flows in particular social settings [33]. The present proposal uses that intellectual precedent to distinguish at least three evidential audiences: the public, affected participants, and authorized independent reviewers. Each audience may have a different legitimate claim to particular information. A defensible arrangement could publish aggregate criteria and reasons, provide claimants with access to relevant individual records, and permit restricted examination of protected material by a competent reviewer. These differentiated channels require justification according to the governing institutional and legal context. They remain vulnerable to risks of collusion, coercion, unauthorized disclosure, and effectively inaccessible appeal.

An important counterexample arises where disclosure itself can expose vulnerable claimants to reprisals. A procedurally open registry may increase public inspectability while reducing safe participation, thereby changing the population whose accounts become available. A seemingly complete dataset would then represent a selected population shaped by disclosure risks. Conversely, a confidentiality rule could be used to shield routine classifications from legitimate scrutiny. Each arrangement calls for evidence about actual access, safeguards, and consequences. The earlier mathematical preservation criteria identify recoverability under specified observation maps; they do not independently authorize a public or protected information channel.

The manuscript therefore treats *protected review* as a conditional institutional option. Its adequacy depends on the reviewer’s competence, independence, permitted information use, record retention, claimant participation, and capacity to obtain consequential action. The mere existence of a confidential review office is insufficient evidence of effective remedy. Context-specific law and professional obligations may also assign priorities that the present conceptual model leaves open.

9.3 Authority, excluded claimants, and downstream decisions

The availability of correction procedures can conceal a second division of authority: the institution empowered to amend a representation may lack jurisdiction to revisit the decisions that relied on it. Section 7 represented a versioned correction and its possible propagation through a dependency system. The reachability result is conditional on the system’s declared links and functional behavior. It does not imply that a legal right of appeal follows from a graph path, nor that a competent decision-maker must alter a determination upon receiving corrected evidence.

A constructed procurement example illustrates the distinction. A regional record custodian corrects the classification of a small community supplier. A separate awarding body has already completed the selection process under its governing rules. The supplier may have a justified interest in explaining the original error, while other bidders may have interests in procedural finality and equal treatment. A model-audit report can document the corrected record, its release version, the applications to which the original classification was supplied, and the applicable review channels. Whether the award may be reopened depends on the responsible institution’s powers and governing norms, including restrictions designed to protect other participants. The existence of the error alone does not determine the legally authorized or ethically appropriate outcome.

A more difficult case concerns an affected claimant excluded by the original domain definition. The record custodian’s accuracy controls may be excellent for admitted submissions, yet the institutional vocabulary may offer no category for the excluded person’s claim. A routine error-correction interface would then be inadequate to address the challenge because the claimant disputes the rule defining admissibility. A relevant review would have to encompass the classification scheme, the evidence of exclusion, and the authority to amend the conditions of admission. This corresponds to a broader concern in epistemic-justice scholarship: recognized channels for testimony and interpretation can structure whose knowledge becomes institutionally actionable [16, 11, 30]. The cited works supply conceptual grounds for examining these mechanisms; they do not verify their presence in the stipulated case.

These examples suggest a three-level diagnostic: correction of a recorded fact, revision of the admission or classification policy, and reconsideration of consequential decisions. Research should trace which authority controls each level and whether affected subjects can approach the appropriate decision forum. The evidence may also show that an existing legal appeal or professional review procedure already addresses the problem adequately. In such settings, the proposed representational audit could add administrative burdens without contributing a distinct safeguard.

9.4 Incomplete histories and irreversible burdens

Historical interpretation presents further limits because institutional records may preserve different forms of ordering and different accounts of responsibility. A timestamp documents a recording convention; a provenance relation documents an asserted relation among records or activities; a causal explanation requires additional evidence; and a normative appraisal requires defensible evaluative premises. Section 5 distinguished these representational functions. They may be correlated without being interchangeable.

Imagine a hypothetical environmental programme that combines reports from several administrative jurisdictions into one annual exposure index. The index is reproducible and appropriate for some comparisons of measured exposure. It may nevertheless omit an event's location within a sequence of decisions, changes in the identity of the competent institution, and the contribution of cumulative burdens to ecological processes that later become difficult to reverse. A subsequent reconstruction could reveal a previously unrecorded exposure or challenge the assumed precedence relation. Such revision would change the admissible historical domain and possibly the questions for which the index was intended. Preservation certified over the original records would remain valid only for that declared domain.

The ethical stakes are heightened where the available evidence is incomplete and some effects may be irreversible. The IPCC documents contexts involving irreversible losses and limits to adaptation [22]; its assessment does not establish the imagined exposure sequence. A responsible institutional inquiry would therefore indicate which ecological observations are available, how omissions affect inference, which uncertainties are material to an intended action, and what independent ecological or community records could challenge the summary. It would also distinguish responsibilities arising under existing rules from substantive claims for recognition, repair, or changed environmental stewardship.

The relational formalism employed throughout the manuscript accommodates multiple historically interpreted orderings and partial records. It does not assume that a global chronological parameter is ontologically fundamental or that a represented possibility is an unrealized event waiting to occur. The relevant question concerns the preservation of explicitly described historical and modal distinctions within a declared analytical task. A new record, a revised interpretation, or the transformation of institutional capacities may alter the space of questions available for inquiry. This is a limit on confidence in a fixed representation rather than evidence that every reduction is ethically defective.

9.5 Comparative stress tests and evidential requirements

A research programme requires observations that discriminate a proposed representational mechanism from familiar institutional alternatives. The proposed account would contribute relatively little if ordinary record correction, a conventional administrative appeal, or substantive resource provision fully explained improved outcomes. Likewise, a model's technical inaccuracy may be sufficient to explain a mistaken decision without invoking epistemic exclusion or improper authority. The present analysis suggests treating these possibilities as competing explanations rather than as variations of the same theory.

Table 19 summarizes a set of constructed stress tests. Each row identifies a particular vulnerability, documentary evidence for investigating it, and an inference that the evidence alone would leave unresolved.

The comparative unit should be an *institutional representational episode*: a bounded inquiry involving an identified representational purpose, an admissible evidence domain, an interpretation or question set, an accountable use decision, and any subsequent challenge or modification. The word *episode* indicates the boundaries of observation rather than a presupposed universal ordering of events. Documenta-

Table 19: Constructed stress tests for the adequacy and governance of reduced representations.

Hard case	Evidence needed	Boundary of inference
Conflicting interpretations	Declared question sets, reasons for admission, dissent records, relevant institutional mandates	Disagreement alone does not identify a uniquely justified representation.
Protected claims	Access rules, disclosure risks, reviewer authority, review outcomes	Withheld public detail does not establish either justified confidentiality or wrongful secrecy.
Excluded claimant	Admission rules, rejected submissions, alternative testimony, competence to revise rules	Missing claims require evidence of exclusion mechanisms and applicable standing.
Corrected record	Version history, dependency logs, reasons for consequential review, governing remedies	Propagated data changes do not ensure changed decisions or effective remedy.
Historical and ecological loss	Provenance, competing event orders, ecological observations, uncertainty assessments	Documentation of precedence alone does not prove causal responsibility.
Claimed audit benefit	Comparator cases, staffing, legal safeguards, selected populations, measured outcomes	Correlation between audit procedures and outcomes does not isolate an audit effect.

tion could include policy versions, dataset specifications, transcripts of admissible hearings, provenance records, protected reviewer logs, allocation decisions, and explanations of unresolved claims. Ethical oversight and information protection would condition access to sensitive materials.

A comparative design would examine episodes with different representational safeguards while attending to institutional competence, available resources, claim-selection processes, legal pathways, and the substantive domain of decision. Researchers could assess whether a declared normative question was recoverable from the published representation, whether legitimate restricted access remedied that information loss, and whether corrected records reached authorized decision-makers. They would then investigate changes in substantive decisions and in affected participants' capacities, using independent evidence and plausible comparators where possible. Model adequacy, access to review, institutional action, and ecological or social outcomes would be measured separately.

Three negative findings would materially qualify the manuscript's programme. First, a declared high-priority query may prove ethically indefensible under closer scrutiny, even though its mathematical preservation is straightforward. Second, existing institutional procedures may provide effective safeguards at lower cost than the proposed representational interface. Third, an apparently improved outcome may result from additional resources, enforceable rights, or changed institutional incentives rather than the representation's refinement. More generally, methodological value would be reduced if the framework systematically failed to identify consequential distinctions that established qualitative inquiry or professional auditing already revealed. These outcomes would motivate revising the interface, narrowing its applicability, or setting it aside for the cases concerned.

9.6 Formal Dependency Audit and the Limits of Integrated Accountability

The formal expressions in this manuscript apply to different questions and often to different domains. Their logical compatibility should be established before treating them as one institutional methodology. The following audit identifies the object, warranted conclusion, and additional inquiry associated with the principal constructions. Its objective is to prevent a chain of locally correct mathematical claims from being presented as an independently justified system of institutional governance.

The representational-policy tuple in Equation 1 identifies decisions that an investigator can record, but it is neither a complete ontology of institutions nor a standard defining their lawful powers. The fiber-wise and family-preservation criteria in Equations 2 and 6 address exact recovery of stipulated answers

Table 20: Formal results, evidential requirements, and remaining normative questions in Paper IV.

Formal object	Evidence needed for a practical claim	Additional ethical or institutional warrant
Question preservation	Justified history domain, answers, coding and interpretation	Authority to choose questions; recognition of excluded claims
Approximate recovery	Valid measurement, error consequences, appropriate loss function	Acceptable burdens and thresholds for a particular use
Historical and epistemic loss	Provenance, competing accounts, admissible external evidence	Standing, confidentiality, competence and credibility assessment
Version propagation	Actual dependencies, corrected records, consequential versions	Reopening powers, procedural finality and effective remedy
Composed institutional pipeline	Stage contracts, collision tests, supplemental evidence channels	Permissible access, authorization and justification of reliance
Institutional feedback	Changed policies, incentives, reporting practices and independent ascertainment	Causal interpretation of altered records and selection processes
Distributed answerability	Activity, role mandate, forum powers and evidence of actual notice	Responsibility, fault, and respective duties of participating bodies
Handoff compatibility	Typed referrals, authorization and versioned receipt records	Competent final review and permitted implementation
Remedial priorities	Feasibility screen, protected interests, uncertainty, resource effects	Substantive priority among permissible and potentially incomparable actions
Evaluative-frame revision	Scope challenges, off-register inquiry, authority to amend rules	Standing, justified closure and legitimacy of revision processes

on declared domains. The approximation bound in Equation 8 applies to a finite domain with real-valued answers and an absolute-error loss. The disclosure proposition in Section 4 concerns incompatible informational demands under explicit observational assumptions. The historical and epistemic propositions use indistinguishability to delimit what can be inferred about particular questions. The reachability result in Section 7 depends on a specified version-dependency structure. The accountability-threshold result in Section 8 characterizes whether a proposed summary can reproduce a stipulated safeguard assignment. None of these mathematical statements establishes a generally applicable priority ordering among rights, costs, privacy, repair, or institutional finality.

The new composition result, Proposition 8.2, connects several of these observations. It supports one inference: once question-relevant distinctions have been collapsed by a deterministic stage, subsequent maps receiving only that stage's output cannot recover them. Equation 14 shows the complementary possibility of question-relative recovery through independent evidence. This distinction avoids a misleading universal description of reduction as irreversible. The irrecoverability claim concerns the specified information channel; it does not prohibit fresh investigation, re-admission of testimony, or reconstruction from independently preserved records. It also does not infer a causal mechanism for a particular outcome from the existence of a record dependency.

Table 20 summarizes the separate warrants required when a formal finding is used in institutional research. The entries are proposed points of inquiry. Whether the resulting evidential record justifies a criticism or remedy depends on the applicable domain-specific substantive standards.

Comparative evaluation should also treat this formal interface as a contestable representational practice. Established professional, qualitative, or legal review may supply equivalent safeguards with fewer burdens; the present notation is valuable only insofar as it improves justified findings or consequential review under specified conditions.

The expanded audit is consolidated into a cross-stage warrant register in Section 9.19. The audit in

Table 21: Epistemic status of the principal result types and potential defeaters in institutional application.

Result type	Claim established under declared assumptions	Evidence capable of changing an applied assessment
Factorization criterion	Specified answers or safeguard assignments are recoverable from a specified output.	A revised question, domain, interpretation or newly warranted policy rule.
Channel-specific collision	A collapsed distinction cannot be reconstructed from that deterministic output alone.	Lawfully obtainable independent records, protected testimony or renewed ascertainment.
Documentary non-identification	A given record is compatible with different competence, recognition or consequence configurations.	Independent instruments, protected reviewer findings or additional historical evidence.
Guarded-path reachability	Specified transitions do or do not admit a compatible route for the specified claim.	Independently authorized fora, lawful referrals or changed jurisdictional rules.
Partial remedial ordering	Certain options remain incomparable under stipulated considerations and feasibility constraints.	Additional rights, evidence, duties or an independently justified priority principle.

Table 20 does not resolve the philosophical foundations of normative claims across scales. That inquiry belongs centrally to Paper III. Paper IV evaluates the conduct and accountability of those who select, maintain and deploy particular descriptions. This boundary is important in the presence of heterogeneous values: mathematical equivalence on a declared query family is compatible with enduring disputes about whose queries warrant recognition and which institutions may decide. The present paper’s proposal is to expose those disputes to better-specified evidence and reasons, retaining the possibility that existing frameworks or institutional practices supply a better answer.

9.7 Proof Status, Defeaters, and Inferential Restraint

The propositions in this paper have unequal inferential roles. The preservation and safeguard-assignment results are *factorization criteria*: given a domain, a reduction and a stipulated answer or policy function, they characterize when the output can reproduce the latter exactly. The finite collision and closed-gate constructions are *conditional impossibility results*: they identify information or paths unavailable within a specified channel or transition system. The documentary indistinguishability results are *non-identification statements*: multiple admissible configurations produce the same observed record while differing on the object of investigation. These are elementary mathematical claims. Their importance in this paper concerns the specific epistemic questions to which they are applied, rather than priority for the mathematics.

A useful proof audit consequently asks three questions for each proposed application. First, is the formal premise established for the particular institution or merely stipulated in a constructed example? Second, is the target of the result a mathematical answer, an institutional power, an observed consequence, or a substantive normative conclusion? Third, which independently available observation could change the finding? A recovery failure may disappear when authorized supplemental evidence is admitted. A handoff gap may be bypassed through an independent competent forum. An internally excluded scope challenge may receive consideration through a legally distinct review channel. Such evidence changes the modeled information or authority structure, without contradicting the original conditional result.

This distinction also specifies what a counterexample can establish. A hypothetical pair of cases with identical published indicators and different institutional powers refutes an inference from the indicator alone to remedial competence; it does not estimate the prevalence of remedial disconnection. A finite successful replay checks that the coded construction has the stated properties; it does not verify that the assumed state space, interpretation or ethical premise fits an actual institution. The comparison protocol must therefore distinguish *formal verification*, *evidential corroboration*, *causal identification*, and *normative assessment*. A conclusion that crosses any of these boundaries needs a stated new warrant. Table 21 is an editorial summary of that inferential discipline, rather than an additional theorem or novel

proof technique.

9.8 Proposition Necessity and the Scope of Formal Originality

The fourteen propositions serve different editorial purposes. They should not be read as fourteen independent mathematical innovations. A result earns its place in this manuscript when it isolates a distinct institutional inference, fixes assumptions needed for a later comparison, or prevents an otherwise plausible but invalid transition from a record to a judgment. The factorization and collision claims deliberately recur across domains because the objects being conflated differ: admission, disclosed description, safeguard assignment, temporal ordering, counterfactual consequences, and completeness of recognition. Their proofs often instantiate familiar set-theoretic or graph-theoretic observations. Their relevance must be assessed through the different practical questions posed, rather than by counting propositions.

Table 22 records the editorial disposition of every proposition. “Retain” signifies a useful distinction within this manuscript, subject to author judgment and the stated assumptions. “Candidate for consolidation” identifies mathematical redundancy that a journal-length version could remove while retaining the argument in prose or a shared lemma. These are editorial assessments, not claims of formal or empirical independent verification.

Table 22: Necessity audit of the fourteen conditional propositions. The grouping concerns manuscript function rather than mathematical originality.

Propositions	Distinct question clarified	Editorial disposition
3.1, 6.1, 8.2	Which admitted or reduced information channel can recover a question-specific claim?	Preserve one general collision criterion; retain stage-specific applications with explicit cross-references.
4.1, 4.2	What approximation bound or disclosure restriction limits a declared question?	Retain: quantitative and access-policy assumptions are different.
5.1, 7.1	Does erased ordering or a missing dependency path defeat a specific recovery or propagation?	Retain as distinct conditional diagnostics; avoid inference of actual injustice.
8.1, 8.3	Can a scalar threshold reproduce a safeguard, or a single observed policy identify a counterfactual consequence?	Retain: policy factorization and causal identification have different targets.
8.4, 8.5	Does a record establish remedial authority, and do local referrals compose into a claim route?	Retain: the first concerns identification, the second relational compatibility.
8.6	Do stipulated feasibility and value dimensions select one response?	Retain as a counterexample to forced total ranking, not a theory of ethical priority.
8.7, 8.8	Can accepted-case records identify excluded persons, and can an internally barred challenge initiate its own review?	Retain: detection completeness and recursive admissibility are separate.

A shorter submission could state the general factorization criterion once and replace some repeated impossibility propositions with corollaries or examples. Conversely, the institutional handoff and recursive-frame constructions require especially careful definition of the applicable state space and authorized transitions, since errors there would change what the model is capable of identifying. The proofs do not warrant claims about population prevalence, causation, rights, institutional fault, or remedies. Those claims require independent records, comparative designs, competent interpretations, and substantive normative reasoning. The editorial objective is not maximal formal density; it is to keep every formal statement traceable to a precise question whose answer could alter an institutional assessment.

9.9 Intellectual Precedents and Incremental Research Contribution

The proposed accountability interface draws together intellectual resources whose central insights are already established in scholarship on institutional classification, epistemic injustice, sociotechnical modeling, statistical reduction, and responsible governance. The aim of this subsection is to specify the narrower contribution that the present manuscript can plausibly make and the comparative evidence needed to evaluate it. A literature comparison supports a statement about conceptual relationships and research questions; it does not establish priority over the complete body of work in these fields.

Bowker and Star show how classifications and standards become consequential elements of information infrastructures [7]. Their analysis establishes a substantial precedent for examining how a category distributes visibility and organizes institutional practice. Fricker identifies testimonial and hermeneutical injustice as distinct wrongs affecting subjects in their epistemic capacities [16]; Dotson and Medina extend attention to silencing, interpretation, resistance, and structured epistemic participation [11, 30]. These works furnish the conceptual foundations for investigating the ethical significance of admission rules and credibility assessments. The present paper's institutional examples do not provide evidence that these wrongs occurred in any particular organization.

Selbst and colleagues investigate failures that arise when abstraction detaches technical assessments from their sociotechnical settings [43]. Their work provides an especially close predecessor for the argument that a formally correct representation may be inadequate for the decision context in which it is used. Gebru and colleagues propose systematic dataset documentation, including motivation, composition, collection, and recommended uses [18]. The W3C provenance recommendations specify interoperable descriptions of entities, activities, agents, and derivations [46, 47]. These approaches already supply substantial means of documenting the construction and circulation of representations. The proposed interface consequently claims no novelty for maintaining dataset documentation, provenance records, or institutional classifications in themselves.

Research on accountable algorithms addresses technical and procedural means of examining decisions, while organizational auditing literature provides procedures for documenting assessments throughout the development and deployment lifecycle [26, 40]. Related institutional guidance includes the NIST AI Risk Management Framework and the OECD AI Principles [31, 36]. Their remit, intended audience, and legal force differ. These sources neither establish universal requirements for every institutional representation nor independently validate the present manuscript's proposed procedure. Jasanoff's account of technologies of humility and Saltelli and colleagues' recommendations for socially responsible modeling additionally emphasize reflection on framing, uncertainty, participation, and consequential public use [23, 42].

The proposed contribution is a *question-indexed, cross-stage accountability test for representational consequences*. For an identified institutional episode, the test locates an admission or classification decision, identifies the associated loss of question-relevant information, traces dependencies among subsequent records and uses, and records the independent or protected evidence by which a competent reviewer could reconsider a consequential determination. This sequence remains conditional on a declared domain, interpretive questions, reliable evidence, actual dependencies, and applicable access and competence rules. It supplies a traceable question for empirical study, rather than a presumption that correction produces a changed outcome.

Table 23 separates the established precedents from the limited additional question proposed here. A claimed contribution beyond these precedents would require evidence that the joint test changes an investigator's warranted findings or an institution's reasoned treatment of a claim. Merely expressing established practices through mathematical notation would provide no such evidence.

A plausible empirical comparison would use the same bounded set of representational episodes under three audit arrangements. One arrangement would assess final model accuracy and aggregate outputs. A second would additionally apply a strong existing practice of dataset documentation, provenance inspection, participatory inquiry, and organizational process review. A third would add the question-indexed cross-stage test. Comparison with the second arrangement is essential: improvement relative only to output accuracy could reflect methods already established in the cited literature. For each arrangement, re-

Table 23: Intellectual precedents and the proposed incremental scope of Paper IV.

Established literature	Established analytical resource	Additional question investigated here
Classification studies and epistemic injustice	Institutional categories, credibility, interpretive exclusion and silencing	Which admission or coding decision causes a particular claim to become unanswerable within a specified evidence channel?
Sociotechnical abstraction and dataset documentation	Context-sensitive critique, documented provenance and collection practices	Can the record distinguish a defect in the admitted domain from information lost at a later reduction stage?
Statistical information comparison and formal reduction	Recoverability or adequacy relative to specified experiments, maps or queries	Which declared question is lost across a particular sequence, and which supplementary evidence could restore review access?
Accountable algorithms and lifecycle auditing	Organizational oversight, audit artifacts, technical inspection and procedural safeguards	Does a correction reach the competent downstream decision process, with a recorded reason for reconsideration or non-reconsideration?
Participatory and responsible modeling	Attention to framing, uncertainty, public involvement and consequential use	Who can challenge the question register and the authorization of a representation for a specific decision?

searchers could record consequential distinctions identified, unsupported accusations avoided, protected claims safely recovered, review decisions adequately explained, and time or resources required for inquiry. These outcomes may conflict, so they should be reported individually rather than compressed into one success score.

The research design would require a reference record assembled through ethically authorized independent inquiry and an appropriate procedure for adjudicating disagreements among reviewers. Otherwise, apparent additional discoveries could be artifacts of unverified allegations or access differences. Blinded or counterbalanced evaluation, where feasible, would help distinguish the effect of the audit method from the expertise and expectations of its users. Comparisons across institutions would also have to account for different legal powers, budgets, documentation quality, selection into grievance channels, and opportunities for consequential correction. Neither a richer audit trail nor a larger count of recognized claims alone demonstrates improved justice.

At least four outcomes would qualify the proposed contribution. The additional test may recover no consequential information beyond a well-conducted existing audit. It may detect formal distinctions that have no supported bearing on the institution’s authorized question. It may generate review burdens or disclosure risks exceeding its plausible gains. It may identify a potential downstream dependency without obtaining evidence that a consequential decision was affected. Conversely, the test would offer incremental analytical value if an otherwise comparably supported audit repeatedly missed a question-relevant loss at an identifiable earlier stage, while the cross-stage inquiry detected that loss, established its evidential significance, and enabled an authorized, reasoned review. Even that result would support a limited methodological claim; ethical evaluation would still require the substantive considerations developed throughout the paper.

The comparison also clarifies the use of *renormalization* in the manuscript’s provisional title. Technical renormalization concerns specific transformations and structures in physics [48]. The institutional constructions here primarily concern aggregation, projection, information selection, and effective descriptions. Their analogy with multiscale physical methods can motivate questions about retained variables and information, but it supplies neither physical dynamics nor an independent ethical calculus. The philosophical precedent for coarse-grained normative facts in List and Valentini [29] reinforces the need for this distinction. The present research contribution should accordingly be assessed as a bounded ethics of representational practice, with its formal techniques serving diagnostic rather than ontological purposes.

9.10 Claim-specific warrants and cross-stage evidential adjudication

The intellectual-precedent comparison in Section 9.9 identifies a methodological proposal whose usefulness depends on more than the detection of formal information loss. The present subsection specifies the evidence needed to move from a description of a reduction map to a warranted institutional conclusion. It treats that movement as a sequence of defeasible inferences, each exposed to competing explanations. This discipline serves two purposes: it prevents mathematical statements from acquiring unwarranted legal or moral force, and it makes the proposed cross-stage method susceptible to meaningful empirical criticism.

A research claim should first identify its *target question*, the *admitted historical domain*, the *interpretive convention*, and the *representation version* being evaluated. Without these declarations, the description “relevant information was lost” leaves the reference class indeterminate. A change in domain can introduce evidence previously outside the sample; a change in interpretation can make a recorded distinction consequential; a change in governing rules can alter the competent institutional question. These changes are analytically different from a technical failure to preserve information under a fixed map. The distinctions developed in Sections 3 and 4 provide the local basis for this first stage.

A *representational warrant* concerns the relation between a specified query and its available descriptions. For deterministic maps, a fiberwise counterexample shows that the query cannot be recovered from the stipulated output alone. The result concerns the functions and domain actually specified. It leaves open whether the query is substantively relevant, whether its inputs correspond to reliable observations, and whether a legitimate reviewer has access to other information. When the relevant domain is contested, researchers can report conditional results for multiple admissible domains and interpretations. The proof remains exact within each declared domain even when its application is uncertain.

An *evidential warrant* addresses what the institutional record establishes about an episode. Recorded refusal, versioned classification, and a published aggregate can support a reconstruction of how an admitted submission was treated. An absent entry, by itself, may also reflect a non-submission, a retention policy, a classification change, an access restriction, or deficient archiving. Independent testimony, protected complaints, provenance records, archival traces, and contemporaneous decision materials may help discriminate among those possibilities. The W3C provenance model documents relations among entities, activities, and agents, and its ontology provides interoperable expressions for such relations [46, 47]. Neither specification certifies the accuracy, completeness, or ethical sufficiency of a recorded episode. Documentary consistency and testimonial credibility require separate assessment, especially where institutional classifications themselves affect whose testimony is heard [7, 16, 11].

A *causal warrant* addresses whether the identified representational choice affected a later institutional result. A dependency edge may establish that an earlier representation could have been consulted; it does not establish that a corrected version would have changed the decision. Evidence might include model access logs, version histories, minutes of deliberation, relevant decision rules, and a bounded comparison with episodes governed by different informational procedures. The investigator should consider resource allocation rules, statutory constraints, independent expert assessments, pre-existing grievances, and administrative capacity as competing explanations. If the same decision would follow from competent consideration of the corrected evidence, the causal claim about an altered outcome should be narrowed accordingly. A consequential correction may nevertheless improve the reasons made available to the affected subjects, and such procedural effects should be evaluated separately from changed substantive outcomes.

An *institutional warrant* asks which body has authority to receive the new evidence, revise a record, reopen or affirm a decision, or supply a remedy. These competencies can be distributed across institutions and need not coincide with the authority to generate the original aggregate. The analysis should record the source and temporal scope of the relevant institutional rule without assuming a universal power of reconsideration. Research on contesting automated decisions offers conceptual forms of procedural and substantive challenge [24, 26]; the related recommendations in the NIST framework and OECD AI Principles concern AI-system governance in their stated scopes [31, 36]. Application to other institutional representations remains an argued extension rather than an inference from these sources. A reviewer

Table 24: Claim-specific warrants, discriminating evidence, and qualification tests for cross-stage assessment.

Warrant	Evidence or argument required	Qualification or competing account
Representational	Declared domain and question; defined reduction; recoverability test	Different domain, question or independent evidence channel
Evidential	Submission histories, records, classification reasons, provenance and protected testimony	Missing archives, alternative coding, restrictions or uncertain credibility
Causal	Version access, institutional reliance, comparison and alternative explanations	Same authorized decision under corrected evidence
Institutional	Documented competence, review rules, access and decision pathway	Lack of jurisdiction, procedural finality or valid referral boundaries
Normative	Stated obligations, protected interests and competing reasons	Defensible disagreement, confidentiality or disproportionate review burdens

who lacks competence to change an allocation may still have authority to amend the provenance record or refer a claim through an authorized channel.

A *normative warrant* then specifies why the observed arrangement would be objectionable or defensible under the relevant ethical reasoning. The mere presence of unequal outcomes, a disputed classification, or non-recoverability within a public index establishes neither unjust exclusion nor a duty of unrestricted disclosure. Ethical assessment may require identifying standing, the significance of protected interests, the burden of review, confidentiality risks, consequences for other claimants, and reasons offered by competing interpretations. Even a formally valid decision within institutional competence can leave a substantively disputed obligation. Normative disagreement should be described through explicit alternative premises and their consequences, rather than reduced to a technical disagreement about a variable.

Table 24 sets out an operational comparison of these warrants. It also records evidence capable of qualifying each type of inference. A negative case can therefore challenge one stage while leaving others intact. A successfully reconstructed provenance chain may coexist with unsupported causal attribution; a supported causal contribution may coexist with lawful institutional action; a legal opportunity for review may coexist with disagreement over procedural accessibility or substantive justice.

Consider a constructed regional programme using a standardized exposure index for adaptation funding. The public index gives the same score to two areas whose submissions include different unregistered care infrastructures. A mathematical assessment may show that the published score cannot answer a question about care-related vulnerability. A dossier obtained through authorized review may establish that both submissions were received but coded under a category that did not retain this distinction. The programme's allocation committee may then consider the corrected records. Three outcomes should be recorded separately: the correction of a description, the provision of reasoned reconsideration, and the alteration of a funding award. If the award follows from an independently binding legal priority that the correction does not affect, there is evidence of representational information loss without evidence that the omitted distinction caused a different allocation. If the review produces a changed allocation, the investigator must still examine the reason stated for that change and whether additional resources or other concurrent decisions contributed to it. In both cases, ethical appraisal requires separate consideration of standing, care-related interests, and the position of other applicants.

The proposed method also faces an *evidence asymmetry*: institutions may possess the records required to inspect their own decisions, while affected participants possess experiential or external evidence poorly represented in those records. An audit restricted to institutional logs may therefore reproduce the very admission boundary it seeks to investigate. Conversely, unrestricted publication of claim-level records may expose confidential information or discourage participation. An appropriate research protocol can distinguish publicly reproducible conclusions from findings reviewed through authorized confidential channels, indicating the access restriction and the limits it imposes on external verification. The authority to conduct a protected review and the legitimacy of a disclosure arrangement themselves

require institutional and ethical justification.

For a comparative evaluation, researchers could pre-register a bounded set of questions, the permitted evidence channels, outcome definitions, and the conditions under which a purported cross-stage defect would be counted. Independent evaluators could then compare the proposed interface with a mature practice that already incorporates provenance reconstruction, qualitative interviews, and organizational auditing [18, 40, 43]. The design should record false attributions, unsupported remedies proposed, relevant distinctions identified, time and resource costs, review burden, and the quality of institutional reasons. The anticipated value of this manuscript’s approach is an additional traceable connection among such practices; its superiority is an empirical question. A finding of no incremental benefit over the stronger existing audit would directly qualify the contribution. This possible outcome gives the paper a defensible boundary rather than a claim of universal methodological necessity.

9.11 Source-specific boundaries of institutional inference

The cross-stage audit method also depends on the disciplined transfer of claims from scholarly and institutional sources into a new analytical setting. This subsection examines four sources of possible overextension: the substitution of a philosophical diagnosis for an observed episode; the transformation of documentary vocabulary into evidential certification; the treatment of an institutional recommendation as an enforceable entitlement; and the projection of a jurisdiction-specific safeguard onto unrelated decisions. The objective is to make citation scope part of the manuscript’s account of epistemic responsibility, while leaving detailed verification records in the accompanying source audit.

The philosophical literature identifies mechanisms whose occurrence calls for case-specific evidence. Fricker’s distinction between testimonial and hermeneutical injustice concerns, respectively, identity-prejudicial credibility assessments and structural disadvantages in the available interpretive resources [16]. Dotson’s analysis specifies testimonial practices of silencing through failures of hearers to respond to the vulnerabilities of speakers [11]. These accounts support the construction of questions for an institutional inquiry. An omitted database entry, an unsuccessful challenge, or a disputed categorization offers a possible investigative starting point. A finding of either form of injustice additionally requires the relevant social and epistemic conditions to be established through evidence and argument.

Technical specifications and policy instruments have other evidential functions. PROV-O is a W3C Recommendation that represents and interchanges provenance information using an ontology of entities, activities, agents, and relations [47]. Its conformance supplies a vocabulary for describing a recorded history; it does not certify the truth or completeness of the account recorded. The NIST AI Risk Management Framework is explicitly voluntary and addresses the risks of AI systems throughout their lifecycle [31]. The OECD AI Recommendation, as amended in 2024, encourages information that makes it possible for adversely affected persons to challenge AI outputs [36]. These instruments can inform good governance by analogy, although they do not in themselves confer jurisdiction over an unrelated institutional record.

A legal example brings the competence limitation into sharper focus. Article 22 of Regulation (EU) 2016/679 addresses specified forms of solely automated individual decision-making, and Article 22(3) attaches particular safeguards to the exceptions in Article 22(2)(a) and (c) [14]. Kaminski and Urban examine competing models for contesting algorithmic decisions, including rule- and standard-based approaches with procedural and substantive orientations [24]. Their conceptual typology offers a resource for institutional design, while the Regulation requires careful determination of jurisdiction, affected subject, decision type, legal basis and applicable safeguards. Table 25 summarizes how the differing statuses of these sources constrain the inferences drawn in this manuscript.

The relevant source-to-claim relationship is therefore a testable part of the argument. For each consequential statement, investigators can record the exact source passage available, whether it expresses a definition, an empirical observation, an evaluative interpretation, or a legal or organizational prescription, and the additional premise needed for its application. Where only a publisher abstract has been examined, the status should be reported as an abstract-level check. Where primary legal or standards text has been read, the finding should identify the provision and the limits of its application. This proce-

Table 25: Source-specific warrants and additional tests for institutional application.

Source domain	Supported analytical use	Additional warrant required
Epistemic injustice scholarship	Forms of credibility-based and interpretive exclusion	Testimony, practice evidence, alternative accounts and ethical reasoning
Provenance specifications	Interoperable descriptions of represented derivations and activities	Record reliability, population completeness and reviewer access
AI governance guidance	Voluntary or recommended methods for system-risk management and challenge	Institutional applicability, adoption and competent authorization
Specific legal provisions	Safeguards for the decisions and subjects covered by their terms	Jurisdiction, material scope, exceptions and legally competent process
Research on contestation	Typologies and comparisons of procedural and substantive challenge	Evidence that a proposed procedure works in the setting under study

dures constrain the argument without equating successful bibliographic verification with full textual or empirical validation.

9.12 Epistemic and normative limits of the proposed framework

The manuscript’s formal vocabulary isolates facts about information recoverability, declared questions, approximation bounds, and dependency relations under explicit assumptions. Its institutional vocabulary identifies authorization, evidence admission, protected review, contestation, and deployment pathways. Its normative vocabulary raises questions about justification, responsibility, standing, care, and remedy. These domains communicate through interpretation and evidence, yet their warrants remain distinct. An exact factorization result proves a property of specified functions. It establishes neither the reliability of the historical observations nor the legitimacy of the selected question. A valid institutional procedure may warrant a decision within its competence without settling the wider question of justice. An ethically persuasive principle may require empirical conditions absent from a particular case.

Several forms of incompleteness deserve continuing attention. A finite declared query family can omit morally material concerns. New evidence can enlarge the domain of admissible histories, and new institutional practices can change which distinctions are recognized as consequential. Some histories and experiences may resist confident translation into the dominant administrative vocabulary. In such cases, attempts at exhaustive formal representation may themselves consolidate the authority of a narrow descriptive language. The proposed framework offers an audit of identifiable choices and losses, together with reasons for revising them; it offers no guarantee of ethical completeness.

The vocabulary of renormalization has an additional disciplinary limit. A physical renormalization-group analysis involves specified transformations of models or degrees of freedom, their composition, and quantities tracked across scales [48]. The institutional reductions examined here generally lack those physical structures. The term appears in the provisional title to signal the ethical questions invited by scale transformations, while the manuscript’s operative models concern aggregation, coarse-graining, provenance, and contextual use. Any stronger physical or mathematical analogy would require separate specification and demonstration.

Finally, public-interest arguments do not eliminate conflicts among privacy, participation, legal finality, administrative feasibility, and distributive consequences. The allocation of burdens among these values involves substantive ethical and institutional choices. Jasanoff’s account of technologies of humility highlights framing, vulnerability, distribution, and learning in governing uncertain technologies [23]; Saltelli and colleagues emphasize transparency and humility concerning policy-facing models [42]. The present framework is compatible with attention to these concerns but does not inherit a validated operational protocol from either source.

Table 26: Research objects and conditions for interpreting results across the four companion manuscripts.

Manuscript and object	Potential interface with Paper IV	Additional warrant required
Paper I: asymmetric generative capacities	A reduced description may obscure the recorded conditions of capacity erosion or restoration.	Historical and causal evidence for a pathway affecting capacity, with alternatives to the proposed mechanism.
Paper II: credible commitments and remedy	Versioned corrections may need to reach a competent reviewer and a consequential decision process.	Established standing, applicable authority, practicable access, review, implementation, and evidence of any resulting change.
Paper III: preservation of normative claims	A question-relative preservation test can reveal which admitted distinctions survive a reduction.	Justification of the chosen query and domain, accuracy of the underlying record, and interpretation of contested claims.
Paper IV: representational accountability	An audit can trace admission, reduction, dissemination, authorized use, and subsequent correction.	Comparison with competent existing audit practice and evidence of additional useful findings, costs, and harms.

9.13 Cross-paper interfaces and their evidential limits

The present section audits the relation between Paper IV and the three companion manuscripts. The purpose is to specify which analytical questions can be transferred between them and which inferences require an additional warrant. The inquiry is a consistency check across provisional research objects, rather than a synthesis that assumes common empirical validity. Paper I considers a conditional mechanism of generative diminishment within asymmetric relations. Paper II examines commitments, contestation, and remedy through institutional pathways. Paper III asks whether selected normative questions can be answered from reduced descriptions of admitted histories. Paper IV examines responsibilities for selecting those descriptions, authorizing their use, and responding to consequential omissions. These objects can be investigated together, but their conclusions occupy different inferential levels.

A shared vocabulary requires a disciplined interface. A *generative capacity* in Paper I concerns the possibilities for continued contribution or development under historically constituted relations; it is not inferred directly from a model's output. *Contestability* in Paper II concerns accessible and competent institutional procedures and their implementation, beyond a statement that a representation is revisable. A *normative claim* in Paper III is evaluated through a declared interpretive question over an admitted historical domain. A *representational decision* in Paper IV includes the choices that establish that domain, the permitted evidence and questions, and the consequential uses of a reduced record. These distinctions define the scope of Table 26.

The interfaces are informative because conclusions from one paper do not automatically establish those of another. Suppose a formal preservation test identifies two stipulated histories assigned the same reduced record, although a declared query distinguishes them. Within its domain the test establishes a failure of exact recovery. It does not by itself establish that the omitted query concerns a legitimate entitlement, that an institution could have collected the needed evidence, or that its omission caused a harmful outcome. Paper IV requires investigation of the representational choices and accessible alternatives. Paper II becomes relevant where a competent procedure can consider a challenge, and Paper I only where the conditions of value circulation and generative capacity are also at issue. The resulting inquiry remains composite: four distinguishable claims need four distinguishable forms of support.

The reverse direction is equally important. A legally competent reviewer might authorize the continuation of a model after a reasoned challenge. That authorization can establish a fact about a particular decision procedure. It does not prove that the retained representation preserves each defensible normative question, nor that the underlying arrangement provides a just distribution of generative opportunities. A correct aggregate may leave a separate historical claim unresolved. Conversely, an independently corroborated complaint can be normatively significant even where it does not supply a complete causal

explanation for aggregate social or ecological change. The cross-paper interface should record these residual questions explicitly.

Constructed case: ecological loss within a shared institutional record. Consider a hypothetical federation that records cooperative returns and environmental obligations in one reporting system. Its published indicator shows stable aggregate productive capacity, while a protected inquiry records localized deterioration in access to a shared wetland. The first task is to determine whether that report and the localized account describe the same admitted historical domain and whether their categories support the ecological question under review. If two histories have the same public indicator and different answers to the declared ecological question, the Paper III preservation criterion diagnoses the informational limitation of that indicator. Paper IV then investigates who chose the categories, what evidence was admitted or excluded, what alternative records are available, and which authority can receive a protected correction. Paper II would examine any review pathway and actual implementation. Paper I would motivate a separate investigation of whether institutional rules contributed to the diminution of regenerative capacities. None of these steps substitutes for the others: the observed aggregate trend, a formal counterexample, procedural accessibility, and the existence of a generative harm are different findings.

This hypothetical case also reveals possible negative results. Additional ecological records might establish that the initial report was explicitly limited to economic production and had never been authorized to adjudicate ecological obligations. A competent existing audit might already document and address the localized loss. The supposed capacity erosion might be due primarily to an environmental event unrelated to the federation's policies. Under those conditions, the representational investigation could remain useful as a clarification of scope while offering little incremental explanation of responsibility. The research should permit that conclusion.

Constructed case: protected evidence and public abstraction. Consider a second hypothetical institution in which the detailed record required to examine a claimant's standing contains confidential testimony. A public aggregation may deliberately omit the identifying features necessary for the standing query. The reduced public record is insufficient for exact reconstruction, as defined relative to that query and its admitted histories. Yet this finding supplies no general reason to publish the protected testimony. The relevant institutional question is whether an appropriately competent and independent reviewer can obtain or verify the evidence under lawful safeguards, communicate a reasoned outcome, and ensure any justified correction reaches the consequential decision process. Paper III clarifies the audience- and query-relative adequacy question; Paper IV investigates the disclosure and authorization choices; Paper II examines the effective contestation route. The privacy obligation and the substantive claim require independent justification. This is a case in which maintaining an incomplete public description may be part of a defensible representational arrangement.

These examples lead to a practical rule for interpreting the research programme. A cross-paper claim should identify the record on which it rests, its historical and interpretive scope, the institutional competence involved, and the kind of conclusion asserted. For instance, a claim that *a correction improved generative conditions* would need evidence of the correction, its transmission and authorized use, and a credible comparison of the conditions with relevant alternative explanations. A claim that *a question was not recoverable from a specified reduction* would need the defined functions, admitted domain, and an appropriate loss witness. A claim that *a remedy was owed* would require substantive legal or ethical premises whose authority cannot be supplied by the loss witness alone. The methodological advantage sought in Paper IV is the ability to keep these claims connected without making them interchangeable.

A comparative audit could therefore test two separate outcomes. The first concerns *discriminating information*: whether the additional cross-stage inquiry identifies consequential distinctions overlooked by a competent audit using established provenance, qualitative and institutional methods. The second concerns *institutional uptake*: whether appropriately authorized actors can evaluate those distinctions and provide reasoned dispositions while observing confidentiality, administrative feasibility, and rights of affected parties. These outcomes should remain distinct. A demonstration of improved detection estab-

lishes no general entitlement to correction, and a completed review establishes no general improvement in justice or regenerative capacities. Appendix B specifies a prospective comparison that may evaluate these differences; its results remain unobserved.

Appendix C operationalizes these distinctions through a structured transfer record and a terminology concordance. The proposed cross-paper consistency criterion is accordingly modest: conclusions should retain their original domains and warrant types when they are used in another manuscript. Formal relations among representations can support diagnostic questions, normative theories can justify why certain claims deserve attention, and institutional studies can assess the actual operation of review and remedy. Their joint application can enrich an investigation when each transfer is independently documented. Where the transfers fail, the failure itself should be reported as a substantive boundary of the programme, rather than concealed within an apparent theoretical synthesis.

The differential evidence register in Section 6.7 adds a qualification to cross-stage review. A discrepancy in published outputs may arise from protected withholding, procedural exclusion, technical non-observation, or other causes; each requires a distinguishable evidential basis and an independently justified access procedure. The framework accordingly treats safe limits on evidential acquisition as constraints on the research method, while preserving uncertainty in judgments that cannot be responsibly corroborated. This remains a proposed audit discipline, subject to comparison with established institutional and qualitative methods.

The feedback analysis in Section 8.7 places a further limit on static comparisons. The introduction of a measure into consequential use can influence who is registered, how a practice develops, which external burdens remain unmeasured, and what an evaluator can infer about alternative policies. Accordingly, an internally correct baseline description may support a narrower conclusion than an institutional forecast about its continued deployment. A responsible comparative audit must preserve the possibility of genuine improvement, adverse displacement, and inconclusive attribution, rather than interpreting every response to a metric as misconduct. The additional evaluation questions in Appendix B.5 are prospective and have not been tested.

The responsibility-interface inquiry in Section 8.8 introduces a further qualification concerning institutional succession and dispersed powers. A documented data custodian, an authorized user of an indicator, and a competent review forum may be different actors. Traceable contribution, institutional answerability, causal influence, and justified remedy must therefore remain separate findings. A proposed typed interface records the stage, actor, forum, evidential basis and scope of competence, while allowing unresolved or protected entries. This interface neither allocates legal liability nor establishes blame. Its incremental value should be assessed through independently adjudicated, protected comparisons against existing public-law, organizational and professional arrangements, including adverse results from erroneous attribution or excessive surveillance. The prospective test in Appendix B.6 has not been implemented.

9.14 Remedial disconnection and the limits of participation

The authority to correct a source record and the authority to reconsider a consequential decision may be divided between institutions. Section 8.9 models this separation without supposing that a successful evidential challenge generates a legal entitlement. One configuration permits a record custodian to correct a classification while leaving the decision maker without reopening powers. Another provides a competent decision maker with no permissible route to the protected evidence. The relevant questions are, respectively, whether information can be examined, whether a competent forum can act on that examination, and which response is justified after competing rights and institutional duties are considered.

Routine case-review statistics can conceal another failure: a person may be absent from the admissible category itself. An investigation should code separately the absence of a submitted complaint, unsafe or prohibitively burdensome access, a jurisdictional rejection, a merits decision without alteration, and an effective reconsideration. These categories imply different defects and different possible remedies. Negative cases matter equally. If established review already links sound evidence with a competent

body and produces reasoned dispositions, an added audit layer may confer little value or impose unnecessary disclosure and administrative costs. The proposed comparison in Appendix B.7 treats the legal and ethical grounds of any reopening as questions external to the formal reachability relation.

9.15 Contestability at institutional boundaries and warranted closure

Continuity of a claim is not equivalent to indefinite preservation of every personal record. Section 8.10 distinguishes a documentary chain, an admissible evidential pathway, a competent review forum and an authorized disposition. A complete archive may coexist with a jurisdictional boundary, while an independent lawful procedure may remain usable after the originating database has been deleted. For a specific challenge, the study must document referral permissions, accessible alternatives, evidential provenance, effective notice, and the reason for any closure.

The strongest tests concern changes in codebooks, service contracts or reviewing authorities for which procedural discontinuity is plausible. A missing handoff should not be inferred solely from absence of a recorded complaint: the claimant may have been unable to access the forum. Conversely, an unbroken version history does not create remedial competence. Comparison with already effective independent review, including its confidentiality and finality protections, can show that an added lineage register is redundant. The proposed protocol in Appendix B.8 preserves uncertainty about unobserved claims and differentiates lawful closure from unresolved exclusion. Actor–forum accountability and participation design are relevant conceptual precedents, although neither supplies case-specific authority [6, 17].

9.16 Plural harms, procedural priority, and remedial restraint

Once relevant distinctions and competent fora have been identified, competing remedies require substantive judgment. Section 8.11 separates urgency of review, admissibility of a proposed response, and its merits in relation to other permissible options. Incomparability under a stipulated partial order does not imply that a decision maker should refrain from acting: applicable law may require timely action, a temporary safeguard, or a reasoned refusal. The model records the conditions under which an ordering is asserted, leaving its moral and legal justification open. Scholarship on plural and potentially incommensurable values motivates caution against replacing heterogeneous interests with one score [44].

Empirical comparison must account for costs to people who do not appear in successful-appeal statistics. A more elaborate procedure may reduce mistaken decisions while increasing disclosure risk, displacing administrative capacity, or delaying protection for other claimants. Researchers should record safe ascertainment, urgent safeguards, resource displacement, justified dispositions and relevant third-party consequences separately. A mature review procedure that achieves equivalent or better results with fewer burdens would limit the proposed interface’s incremental value. Appendix B.9 defines the prospective measures without determining a general priority rule.

9.17 Ethical criteria as governed and revisable objects

A challenge may concern the admission rule used to decide which challenges can be heard. Section 8.12 distinguishes review of an individual classification from review of categories, evidential conditions, participation rules and evaluative priorities. Proposition 8.7 establishes only that an accepted-case register can be compatible with distinct histories of off-register exclusion. A case study therefore needs independent evidence about the potential claimant population, the scope of review competence and the consequences of a revision; the proposition does not select a just jurisdiction or grant standing.

An adequate revision record should indicate the disputed rule, the authorized decision-maker, the reasons for admitting or rejecting the request, treatment of pending claims, and safeguards for protected or absent third parties. More participation is not independently proof of better outcomes, and complete version histories may coexist with substantively inaccessible rule review. Fraser’s attention to contested frames of justice and Anderson’s account of democratic-epistemic learning support the philosophical

inquiry, but their distinct normative commitments are not adopted as a ready-made institutional procedure [15, 3]. Appendix B.10 sets out comparison against competent existing arrangements, including inconclusive and adverse findings.

9.18 Reflexive governance, justified closure and normative transformation

The problem of recursive rule review concerns the institution’s authority to revise the criteria that determine which revisions it considers. Proposition 8.8 is a reachability statement: if every internally reachable rule state excludes a specified meta-challenge, internal procedures alone cannot establish its consideration. The statement offers no proof of the challenge’s normative validity and no prescription for an external forum. A competent institution may reasonably reject, defer or limit a claim, subject to applicable duties and reason-giving requirements. The relevant distinction is between a justified disposition and exclusion that cannot be independently examined.

Revision can alter the evidential population, the costs borne by existing applicants and the practical consequences of disclosure. Assessment must therefore distinguish conferred authority, evidential adequacy and substantive ethical justification, while tracking plausible rival explanations for any observed changes. Independent oversight may already provide the needed route; conversely, new review procedures may compromise safety, delay mandated decisions or merely reproduce the earlier admission boundary. Appendix B.11 treats these cases as prospective comparisons rather than empirical confirmation of the proposed governance model.

9.19 Integrated Warrant Architecture for Representational Accountability

The purpose of this subsection is to consolidate the manuscript’s conditional results into an integrated account of institutional representational practice. Sections 2–7 identify information and review problems at distinct representational stages. Section 8 examines feedback, distribution of competence, institutional succession, remedial prioritization and revision of evaluative frames. Integration requires an account of how those inquiries bear upon one another while respecting their different objects. The resulting *warrant architecture* is a proposed method for organizing investigation. It specifies which transitions require justification, how failure at one transition can affect later inquiry, and what evidence might change the analysis. It is neither an all-purpose moral decision procedure nor a condition sufficient for a legally valid remedy.

An inquiry-indexed pathway and its scope. Let u denote a specified consequential use, q a declared question whose relevance is independently open to challenge, and k an identified institutional configuration. An accountability inquiry identifies the admitted history domain H_k , its representational policy and recorded outputs, admissible supplemental evidence, access arrangements, competent fora, possible dispositions, and the criteria under which those elements may be revised. The configuration is not assumed to be a snapshot of an ontologically given present. It identifies a historically and institutionally interpreted setting in which claims and records may have documented dependencies; comparisons between configurations require an expressly justified ordering or transformation relation.

A useful coding device for a *claim episode* is the typed warrant profile in Equation 28. Here the subscripts identify separate assessments, and the three-valued entries record corroboration, a warranted negative finding, or unresolved evidence, rather than moral approval or disapproval.

$$\mathcal{W}(q, u, k) = (w_{\text{question}}, w_{\text{evidence}}, w_{\text{access}}, w_{\text{competence}}, w_{\text{transfer}}, w_{\text{response}}, w_{\text{frame}}) \in \{1, 0, ?\}^7. \quad (28)$$

The *question* entry asks whether a material question and its domain have been specified and defended for this inquiry. The *evidence* entry concerns the information available for answering it, including independent corroboration and uncertainty. The *access* entry records whether a legally and ethically permitted route to relevant information exists for an authorized reviewer. The *competence* entry records a verified forum and its powers for the specific claim. The *transfer* entry concerns any required propagation across

representational, organizational or decisional dependencies. The *response* entry concerns institutionally available actions and their constraints, while the *frame* entry concerns the potential for a challenge to the rules setting admission, standing or evaluation criteria. Each field must be accompanied by its evidence, interpretation, responsible assessor and relevant restrictions; a numeral alone has little meaning.

This profile is a *disaggregated diagnostic register*. Its components are not commensurable quantities, and no sum, product, threshold or lexicographic ranking is proposed. A zero may indicate a reasoned absence of competence, an absent document, a justified exclusion of access, or a failure of an attempted transfer, depending on the field and record. An unknown value may be appropriate when the governing authority or independently ascertainable claim population cannot be established. The same observation can support competing normative conclusions under different lawful purposes or protected interests. In particular, a positive frame-review entry does not certify democratic legitimacy, and a negative entry does not independently establish injustice.

Noninterchangeable limits across the pathway. Several elementary results in the manuscript acquire a clearer relation when considered in this register. The fiberwise preservation conditions in Section 4 address whether a declared answer is recoverable from a stipulated reduced representation. Proposition 8.2 shows that deterministic composition cannot restore a distinction already collapsed if its only input is the collapsed output. Such a *representational cut* can remain even when every institution handling the later record acts competently. New admissible evidence may change the effective information channel, as Equation 14 makes explicit; the original reduction does not thereby retrospectively retain a distinction it had discarded.

A different, *institutional cut*, arises when relevant evidence can reach an actor who lacks the power to reconsider the consequential use. The responsibility-interface and authority-separation analyses in Section 8 show why a documentary path, an answerability forum and an effective remedial route must be verified separately. Proposition 8.4 supplies a narrowly specified indistinguishability example: equivalent accessible logs can coexist with distinct authorized remedy routes. Proposition 8.5 establishes an additional point about composition: a sequence of nonempty local referral relations need not yield a compatible end-to-end route for a particular claim. Information preservation, available review and effective competence consequently refer to different types of conditions.

A third, *evaluative cut*, occurs when a process for correcting admitted claims cannot consider whether its own admission or revision rule is warranted. Proposition 8.7 limits conclusions that can be drawn from an accepted-case register about independently ascertainable recognition gaps. Proposition 8.8 addresses a narrower reachability question for internal-only rule-review transitions. Neither result determines rightful standing or mandates indefinite reopening. A legally independent forum, a protected procedure, or warranted closure may alter the institutional assessment without altering the validity of the elementary formal observations. Table 27 brings these distinctions into a single documentary map.

A comparative construction. Consider a hypothetical programme that allocates support through a published completion-rate indicator. Two eligible requests can generate the same published success rate, although one history contains an unsuccessful protected appeal and the other does not. An authorized archivist in either configuration can retain the original evidence; a public summary cannot recover that distinction by itself. Suppose further that a reviewing office receives notice of the discrepancy. In configuration k_1 , that office can demand reasons and refer the matter to a decision-maker competent to reconsider the allocation. In configuration k_2 , the same notice is received but the office's powers stop at correcting an archival category, and the referring rule does not authorize the decision-maker to reopen the allocation. The public aggregate and the recorded acknowledgement may be identical; the competent routes differ. Finally, in configuration k_3 , an independent body can hear a challenge to the eligibility criterion itself, subject to protected participation and an authorized decision about pending cases. These are *stipulated* differences, not observations of real administrative conduct or claims that a particular remedy is owed.

The comparison generates distinct research questions. To identify the appeal-related distinction one

Table 27: Integrated cross-stage warrant map. The categories are proposed audit questions, not universally mandatory procedural steps or a ranking of ethical values.

Inquiry register	Illustrative evidence and verification	Question left open by that verification
Question and domain	Purpose, eligible histories, alternative questions, evidence of omitted accounts	Which questions and participants warrant recognition?
Representation and evidence	Policy version, coding record, collision test, protected independent evidence	Is a retained distinction reliable and material to the actual use?
Protected access	Authorized disclosure rules, reviewer access, safety assessment, recorded refusals	Does the information flow respect confidentiality and third-party interests?
Answerability and competence	Delegation, jurisdiction, forum powers, reasons and challenge procedures	Can this forum reconsider the claim or merely acknowledge it?
Handoff and dependence	Compatible versioned referrals, receipt, dependency paths, implementation evidence	Can a competent response reach the consequential decision?
Feasible disposition	Lawful options, material uncertainty, urgency, rights screening and capacity effects	Which action, if any, is substantively justified?
Revision of criteria	Standing rules, scope challenges, revision powers, pending-case treatment	Who can contest the criteria governing these judgments?

needs independent evidence, with a justified permission to inspect it. To establish a right to reconsideration one needs the governing legal or institutional instrument and evidence of the actual competence of the forum. To assess whether a remedy changed consequential use one needs records of reliance, a lawful disposition, and a compatible transmission route. To evaluate the eligibility rule one needs a warrant for admitting the scope challenge and, where possible, independent ascertainment of affected interests beyond the programme's accepted-case log. Equal aggregate accuracy cannot answer these questions. Conversely, discovering the hidden appeal does not demonstrate unlawful conduct or establish that reopening the allocation is warranted. This construction exhibits an *analytical separation* of warrants, without claiming a new theorem or a validated policy intervention.

Institutional feedback and the selection of evidence. The architecture also clarifies the feedback account in Section 8.7. A category used to allocate resources may change how participants classify activity, document need or attempt a challenge. Those practices in turn affect which histories are admitted under a subsequent institutional configuration. Thus a pathway is not adequately described solely by a fixed information-reduction chain when institutional use changes the production of its inputs. The appropriate extension is a comparison of declared configurations and the evidenced transformations between them, with attention to alternative causes. The formal profile in Equation 28 should consequently be versioned together with the institution's query family, categories, powers, disclosure rules and material policy dependencies. A changed profile is evidence of a changed assessment record; it is not itself evidence that the institution has become more ethical.

Care is needed where a correction alters the evidential population. If formerly excluded claims enter the register, an increased number of detected omissions may indicate improved ascertainment rather than deterioration in representational practice. If a new confidentiality rule reduces public case visibility, lower published complaint counts may coexist with effective protected review, with harmful suppression, or with both in different subpopulations. A predeclared denominator and a secure independent ascertainment strategy are therefore necessary for interpretable comparison when feasible. When they are unavailable, investigators should report the resulting identification limit rather than infer an improved or worsened condition from the administrative trend.

Plural values, response authority and the refusal of an ethics score. The final transition from a documented defect to a proposed remedy presents a substantive ethical question. Section 8.11 provides a

constrained comparison of feasible responses under stipulated plural considerations. An accessible complaint route may coexist with legally protected finality; restoring one claimant's record may disclose another person's protected account; a narrow case correction may be available while category-wide reform would require a separate mandate. Where the candidate actions remain incomparable, the documentary profile does not license an arithmetic aggregation of rights, costs, uncertainty and urgency. Decisions may still be required under applicable duties. Their justification must identify the authorized decision-maker, the relevant substantive reasons and the consequences of chosen action or justified refusal.

The ethics of renormalization can therefore be stated more precisely as an inquiry into *which distinctions are made institutionally available for which justified questions, through which authorized evidential and decisional relations, and with what consequences for the continuing ability to recognize and address claims*. This formulation covers the construction of descriptions, their institutional circulation, the generation of further records through consequential use, and the governance of the criteria for review. It retains the earlier distinction between representation and reality: the investigative interface can be improved without identifying the represented record with the full relational history, and without assigning an automatic moral status to greater detail or broader access.

Comparative test and stopping conditions. An empirical assessment should compare a versioned cross-stage warrant register with a mature, competent baseline combining case review, provenance documentation, confidentiality procedures and existing rule-change mechanisms. The primary inferential question is whether the register improves the independent detection and authorized disposition of materially omitted distinctions, taking into account the burdens and adverse effects it produces. Observation of a complete register alone is insufficient. A prospective investigation should report separately the accuracy of coding, detection of off-register claims, lawful reviewer access, verified competence, decision-level transmission, actual dispositions, harms to third parties, resource displacement, and changes attributable to other reforms. Appendix B.12 specifies a bounded extension of the proposed protocol.

The register should be abandoned, simplified or restricted when competent existing processes perform equivalently with fewer burdens, when the required evidence cannot be acquired safely, when the proposed classification imposes material interpretive distortion, or when the operative law and institutional mandate prevent a proposed procedural step. This is a methodological stopping condition, not a universal ethical verdict. The formal account supports a disciplined inventory of questions and evidential dependencies. Evaluating legitimacy, protection, responsibility and justice remains the task of domain-specific inquiry with affected participants and competent institutions.

9.20 Release Claims, Discriminating Tests, and Evidential Boundaries

A public-facing account of this inquiry should distinguish three levels of claim. The manuscript *establishes conditionally* mathematical facts about specified descriptions, observation channels, decision dependencies, and permitted transitions. It *proposes* a cross-stage practice of specifying the question, identifying the earliest relevant information or competence gap, and documenting the warrant needed for review or response. It *does not establish* that this practice improves institutional outcomes, identifies every affected party, creates jurisdiction, or yields a uniquely just remedy. These three claim levels should remain separate in the abstract, any dissemination materials, and derivative applications.

The most discriminating empirical comparison is not whether an institution can fill all seven fields in the integrated profile. It is whether its use reveals an independently corroborated issue that would otherwise remain unrecognized, produces a lawful and justified opportunity for assessment or response, and does so without imposing disproportionate investigative costs, risks, or losses of participation. A positive demonstration would require at least: (i) predeclared questions and eligibility definitions; (ii) independently acquired and ethically usable reference evidence for identified omissions, including plausible off-register claims; (iii) a competent baseline such as an existing legal review, professional audit, or archival investigation; (iv) documented differences in findings, reasons, or dispositions; and (v) assessment of confidentiality breaches, false allegations, administrative burden, and displacement of other

legitimate work. Evidence of complete registers or more recorded complaints alone would fall short.

The proposal could fail for principled reasons. A protected independent forum might already examine exactly the missing distinctions and provide a satisfactory response. A systematic register might reproduce the original institution's taxonomy and suppress a form of testimony that is difficult to encode. It might generate new files with weak evidential provenance, incentivize strategic coding, increase burdens on claimants, or expose protected third parties. It might also correctly detect a difference but have no permitted path to change an allocation. These possibilities are not auxiliary exceptions: each is a potential reason to narrow, redesign, or abandon the proposed procedure in the relevant setting.

A responsible release should also distinguish *bibliographic verification* from *passage-level support*. Checked publisher descriptions, abstracts, metadata, and standards passages help establish what a source is plausibly about; they do not demonstrate that every interpretive sentence in a long manuscript accurately reflects the cited full text. Nor does successful compilation verify proofs, legal applicability, or the comparative protocol. The accompanying audit records known limits rather than converting unfinished checks into a declaration of publication readiness. The author's approval and independent substantive review remain necessary before this manuscript is represented as a finalized scholarly contribution.

9.21 Open problems and concluding synthesis

A prospective comparative protocol for examining the proposed accountability interface is specified in Appendix B. Its design separates observable corrective processes from causal and normative conclusions, and includes negative cases in which existing audit practices perform equally well or better.

The remaining problems fall into four connected research tasks. First, independently ascertain claims beyond the currently admitted histories and questions, including ecological and relational dependencies that resist existing categories. Second, determine when protected access, lawful referral and competent review preserve actionable contestability across institutional changes. Third, study the legitimacy and distributive effects of revising evaluative frames and allocating attention among claims, accounting for third-party rights and the burdens of participation. Fourth, identify the incremental effects of any proposed audit procedure against mature existing practice, including displacement, unsafe disclosure and negative or inconclusive results. Each task requires the information and normative warrants appropriate to its own question.

The integrated account in Section 9.19 provides a common language for these problems without equating a completed register with legitimate action. Its claim is narrower and more useful: once a representation becomes consequential, an investigation should be able to say which declared question is at issue, where a distinction was lost or rendered inaccessible, what additional evidence is warranted, who is competent to consider it, and whether an authorized response can reach the relevant decision. Failure at any one interface warrants a specifically scoped inquiry, not an automatic condemnation; success at each interface is not a proof of justice.

A completed inquiry may support several different dispositions, which should remain analytically distinct. An independently corroborated representation may prove adequate for the declared question while leaving other, unexamined questions open. A demonstrated informational omission may justify record correction and a competent reconsideration procedure. A materially established omission may also exceed the investigating forum's authority, in which case the record should specify the permitted referral, available independent forum, or reason for a justified institutional limit. Finally, an inquiry may end with evidence too weak for a substantive finding, a protected source that cannot appropriately be disclosed, or an off-register population whose relevant claims remain unknown. This last disposition should identify uncertainty without classifying unobserved persons as compliant, unaffected, or ineligible. These are proposed documentary conclusions, not an exhaustive taxonomy of outcomes or a new source of remedial power.

The ethical stakes reside in the practical and historical consequences of making certain accounts actionable while others remain difficult to formulate, preserve, or contest. Those conditions may change as classifications influence conduct and as the authority to revise classifications is itself challenged. A

question-indexed, versioned examination can make some dependencies more inspectable. It remains only one instrument among legal, historical, ethnographic, technical, and participatory forms of judgment. The paper’s bounded conclusion is that the responsible use of multiscale institutional representations requires reasons that connect information choices to legitimate evidential and decisional processes, with openness to justified revision. Whether the proposed integrated procedure supplies better reasons or better outcomes than mature existing practice is an empirical and ethical question left for independent testing.

A A Constructed Cross-Stage Accountability Demonstration

This appendix supplies a small reproducible demonstration of the cross-stage accountability interface developed in Sections 7–8. Its purpose is methodological: to isolate the relations among representational sufficiency, admissible supplementary evidence, a version update, institutional competence, and consequential reconsideration in a finite stipulated case. The accompanying source archive includes a JSON case specification and a Python script using only the standard library. All cases, parameters, decision rules, and outcomes are invented for illustration. They describe neither actual administrative conduct nor observed improvement in justice.

A.1 Declared histories, questions, and information channels

Consider a hypothetical coastal-allocation programme receiving two admissible candidate histories, h_0 and h_1 . Both histories contain the same exposure score, 7, below a stipulated general eligibility threshold, 8. They differ in a protected indicator q describing whether additional care-related vulnerability has been established through an authorized evidential procedure: $q(h_0) = 0$ and $q(h_1) = 1$. The programme publishes only the exposure score. With $H = \{h_0, h_1\}$ and $R_{\text{pub}}(h_i) = 7$, the published mapping identifies these histories even though their answers to the declared question differ. Accordingly, there exists no function f satisfying $q = f \circ R_{\text{pub}}$ on H . This is an application of the fiberwise criterion in Equation 2, rather than a new mathematical result.

The indicator q is stipulated to be available in an independent, protected dossier. Its availability does not settle whether that dossier is reliable, whether the indicated vulnerability should alter allocation priorities, or whether disclosure is lawful. In the constructed example an authorized reviewer may inspect the dossier under an explicitly specified access rule; the public index remains unchanged. The distinction corresponds to the documentary separation between recorded provenance and warranted interpretation discussed in Section 3. PROV-O supplies a vocabulary for exchanging provenance assertions, but it does not independently validate the truth of a dossier or establish its institutional admissibility [47].

A.2 Versions, authority, and decision rules

Five replay conditions are defined for the history h_1 . Each condition specifies whether the originating source has been corrected, whether a notice of the new version reaches the downstream institution, whether the receiving reviewer may inspect the protected dossier, whether a competent authority can reconsider the decision, and which allocation rule has been adopted. In the *flexible* rule, a verified care indicator permits eligibility in addition to the exposure threshold, subject to the stipulated competence of the reviewer. In the *fixed* rule, eligibility depends exclusively on the exposure threshold. The labels “flexible” and “fixed” describe assumed decision functions; neither receives ethical endorsement from the example.

For this demonstration, the initial decision is a denial under the threshold-only information channel. A source correction is recorded separately from notification, protected evidence access, and authorized reconsideration. The downstream decision may change only if each necessary condition in the modeled pathway holds and the applicable rule gives the corrected evidence decisional relevance. Table 28 reports

Table 28: Constructed replay of a representational correction across institutional stages. The outcomes follow the declared toy rules and are not empirical estimates.

Replay	Correction and review pathway	Decision	Interpretive boundary
Baseline	No source correction or protected review	Denied	Public score lacks the declared distinction
Source-only	Source corrected; downstream notice absent	Denied	Local correction has not reached the user of the old version
Notice-only	Correction delivered; receiving body lacks review competence	Denied	A notification supplies no independent reopening authority
Authorized review	Protected dossier admitted; notice received; competent flexible-rule review	Eligible	Outcome changes under the stipulated rule and evidence
Fixed-rule review	Same procedural access; competent reviewer applies binding exposure-only rule	Denied	Review occurs without an altered substantive decision

the script-generated outcomes. The entries use a consistent history, exposure score, and protected vulnerability indicator so that the effects of stipulated institutional differences can be seen without changing the underlying case.

The result in the authorized-review condition is a logical implication of the assumed rule and inputs. It cannot establish the causal effect of an audit in an actual organization. In particular, the fixed-rule case demonstrates that evidential recovery, procedural reconsideration, and substantive outcome modification are distinguishable. The notice-only condition demonstrates that propagation of a corrected record does not confer competence to reopen a decision. A separate legal inquiry would be necessary to determine whether either rule, the access restriction, or the scope of reconsideration is justified under applicable law. The contestation arrangements studied by Kaminski and Urban concern algorithmic decisions and provide a comparative procedural precedent, not a universal rule for the invented programme [24].

A.3 Replication, comparison, and failure conditions

The bundled file `examples/replay_cases.json` specifies the candidate histories and the five replay conditions. The bundled script `examples/replay_accountability.py` evaluates the declared cases, asserts that the public score cannot recover q across the candidate histories, and checks the resulting decisions against expectations. It also verifies that an admitted protected dossier does not automatically change the decision under a binding exposure-only rule. No personal or confidential records are processed. Running the script reproduces the internal logical outcomes, and its tests can reveal an inconsistency between the prose, the source specification, and the stipulated decision functions.

A substantive evaluation of the proposed interface would require a stronger comparator than a public-score-only baseline. An institution may already combine record correction, qualitative review, provenance assessment, and competent appeal. Accordingly, a research design should compare the cross-stage interface with such a mature practice. Investigators would need to measure additional material distinctions identified, mistaken causal attributions, unsupported remedies, administrative and participant burdens, adverse disclosure events, and the quality and practical effects of institutional reasons. The comparison should also evaluate decisions that remain unchanged after reasoned review and claims that cannot safely be disclosed. A finding that an established process offers equivalent or superior outcomes at lower cost would qualify the proposed interface’s incremental utility.

The demonstration therefore establishes only internal reproducibility within an explicit finite specification. Its ethically significant questions remain open: whether the protected indicator deserves decisional priority, how reviewer competence is constituted, whether rejected claims can obtain standing, how evidential errors are corrected, and how the burdens of disclosure and reconsideration should be al-

located. These require independent empirical, legal, historical, and normative inquiry. The worked case is a device for exposing dependencies in an evaluable form, consistent with the manuscript’s bounded conception of representational accountability.

B A Prospective Comparative Evaluation Protocol

The constructed replay in Appendix A establishes internal consistency for a stipulated decision system. This appendix specifies a distinct, prospective strategy for investigating whether the paper’s proposed cross-stage accountability interface adds practical value to existing institutional auditing. Its role is to make the proposed comparison open to critical inspection before cases are selected or outcomes are observed. The protocol is a research design, rather than a report of an implemented study. No institutions, participants, interventions, or evaluation results are claimed to have been recruited, observed, or completed.

B.1 Research object and comparator

The intervention to be evaluated is a documented *cross-stage representation audit* applied to a consequential institutional decision pathway. The intervention records the admissible historical domain and the declared question family; reconstructs the admission and classification stages; identifies the versioned representation consumed by a downstream body; specifies the authority that could receive a correction; and records whether the correction receives an admissible review and a reasoned disposition. This is a documentary and interpretive procedure, supported where appropriate by existing provenance records. It is not a new statutory entitlement, an automatic correction system, or an intervention that independently determines substantive outcomes.

The comparison condition should be a mature, reasonably resourced institutional process for case review, provenance checking, qualitative investigation, and independent escalation where relevant. Comparing an elaborate protocol against a record-keeping system without basic review facilities would exaggerate its incremental benefit. The operative comparison therefore concerns the *additional* distinctions identified, transmitted, and competently evaluated, together with the costs and harms of obtaining them. Organizational auditing scholarship offers relevant precedents for a lifecycle view [40]; the particular comparison remains a proposal for investigation.

The evaluation is concerned with three separate objects: (i) representationally material distinctions, (ii) the successful disposition of a claim through a competent review pathway, and (iii) substantive institutional outcomes. The first may be established by documentary reconstruction; the second requires recorded procedural action; the third additionally requires defensible causal and normative analysis. Counting any of these as an interchangeable measure of ethical improvement would collapse the paper’s central distinctions.

B.2 Study population, inclusion, and sampling

A possible observational unit is a *claim episode*: a request, concern, or challenge related to a consequential representation, together with the institutional history relevant to its handling. Episodes can be nested within institutions, programme domains, and decision pathways. The unit must be fixed before data collection, with rules for repeated filings, merged claims, referrals, and multiple represented subjects. Institutions may differ in competence, resource availability, information infrastructure, and conditions of public access. Such differences warrant stratification or explicit contextual analysis rather than automatic pooling.

The largest selection problem concerns episodes absent from official registers. An administrative dataset conditions observation on prior recognition and intake. A defensible sampling frame should therefore combine, where lawful and feasible, institutional registers with separately collected evidence on unsuccessful submissions, rejected intake attempts, community reports, and archival discrepancies.

The independent channels require informed and voluntary participation, protected handling, and attention to who is able to come forward. These measures improve coverage without guaranteeing that every affected claimant becomes observable. Where the wider population cannot be characterized, findings must be restricted to the observable or independently reconstructed population.

An evaluation should record a flow of sampled episodes through eligibility, evidence availability, admission, review and outcome ascertainment, distinguishing absence of documentation from an affirmative determination that an event did not occur. The protocol does not prescribe a universal number of episodes or institutions; power, precision, recruitment burden, costs, and the heterogeneity of exposure should be established through a context-specific feasibility analysis and a publicly documented rationale before the evaluative analysis.

B.3 Intervention assignment and causal interpretation

When ordinary legal protections and access to remedy can be maintained for everyone, a prospective study may compare a standard audit workflow with the same workflow augmented by the cross-stage interface. For example, the additional documentary audit could be introduced sequentially across comparable teams under an authorized implementation plan. Such an evaluation should preserve standing, legal representation, review competence, and obligatory safeguards under both conditions. Allocation of protections, access to evidence, or ordinary appeal rights as experimental variables would require a different legal and ethical justification and is outside the present proposal.

The causal question must define the intervention, assignment unit, start of exposure, comparison process, follow-up rules, and outcome before analysis. The target-trial formulation developed by Hernán and Robins provides a methodological precedent for clarifying a comparison even when feasible evidence is observational [20]. Where randomized introduction of the supplemental audit is infeasible, candidate alternatives include closely documented matched comparisons, staggered implementation with independently supported identifying assumptions, and structured process tracing. Such designs should report plausible confounding by case severity, staffing, legal entitlements, public enforcement, resource transfers, documentation quality, and changes in claimant composition. An observational association is not sufficient to establish the effect of the proposed interface.

If an institution adopts the augmented protocol precisely because it has experienced an unresolved failure, differences between audited and unaudited cases may reflect this prior failure. Similarly, skilled reviewers may both choose the new protocol and improve case handling independently of it. These mechanisms require attention in analysis and may preclude a causal conclusion. The study should present noncausal descriptive comparisons when the identification assumptions are unpersuasive.

B.4 Endpoints, denominators, and independent adjudication

The proposed primary endpoint is the proportion of pre-specified, independently eligible episodes for which a *material representational omission* is both (a) corroborated by an independent adjudication process and (b) transmitted to a competent review body with a recorded disposition within an explicitly declared observation horizon. An omission is material only relative to a predeclared question and a specified institutional use. The independent assessors should review protected evidence under approved access arrangements, with a procedure for contesting their own classification and recording disagreement. Their adjudication remains fallible and should not be described as an ethical ground truth.

Secondary endpoints should be reported separately: identification of previously unregistered claims; accuracy and timeliness of correction notices; access to competent reconsideration; quality and specificity of reasons; reversals or modifications of decisions; persistence of substantive effects; participant time and financial burdens; administrative costs; error in attribution; and confidentiality or retaliation incidents. Unchanged decisions after meaningful review require a distinct category from cases in which no review occurred. Increased detection of omissions can initially raise recorded failure counts; interpretation must therefore distinguish greater observability from worse institutional performance.

Table 29: Proposed comparative evaluation ledger. Each endpoint requires a separate source and an explicit limitation.

Assessment domain	Required evidence	Principal interpretive limitation
Claim coverage	Registered episodes and independently sourced intake attempts	Unseen claimants and self-selection remain possible
Representational omission	Versioned source, declared question, independent admissible dossier	Corroboration depends on a contestable interpretation
Correction delivery	Audit log, recipient acknowledgement, authority map	Delivery does not imply competence or reconsideration
Reasoned review	Review order, disposition, reasons, opportunity to challenge	Procedural completion does not settle moral legitimacy
Substantive consequences	Decision records, enforcement, resources, affected conditions	Concurrent interventions and confounding may explain changes
Safety and burden	Access logs, adverse-event reports, participant testimony, cost records	Underreporting and differential vulnerability affect measurement

The denominator for each endpoint must be stated. Reporting a success rate only among cases admitted for review can conceal exclusion at intake. Reporting success among all received cases can obscure unrecorded claims. Where the number of potentially affected but unobserved episodes is unknown, an overall coverage proportion cannot be identified from the official register alone. The evaluation should present the defensible observed denominator and, when data permit, transparent sensitivity bounds under declared assumptions about nonregistration and missing outcomes.

Table 29 summarizes evidential distinctions that must remain visible throughout the evaluation.

B.5 Policy reactivity and longitudinal interpretation

A prospective evaluation should additionally identify whether the audit intervention, the announced outcome measure, or a subsequent institutional use could change the production of the observations being assessed. Before collecting outcomes, investigators should register the operative version of each indicator, material incentives attached to it, changes in intake and eligibility definitions, and the locations at which protected or independent sources may provide admissible corroboration. Where organizational arrangements permit, a process study could compare periods or units subject to different forms of indicator use, taking account of concurrent policies, exposure, anticipation, and possible spillovers. Such comparisons would require case-specific identification assumptions; the protocol claims no general treatment effect from a descriptive trajectory alone.

The analysis should separate a rise in actual service quality or accessibility from a rise in coded completion and should preserve information about who enters the measured denominator. Any indicators used to evaluate the audit must themselves be assessed for incentives to selective reporting. Investigators should predeclare possible contrary results: no detectable reactive change; genuine improvement without record distortion; changes attributable to independent administrative reform; and unresolvable confounding. The inclusion of the reactivity module should be evaluated for additional detection, reviewer burden, safety, and cost relative to the comparator. Research participation, secure records access, adverse-event handling, and decisions to modify operational indicators remain subject to competent institutional approval. This module is a proposed safeguard against interpretive overreach, not an implemented longitudinal study.

B.6 Distributed answerability and institutional-succession module

The additional responsibility-interface procedure in Section 8.8 is an unimplemented audit component. A prospective comparison would examine whether this component detects material gaps in institutional answerability *beyond* those identified by a mature existing audit with competent reviewers. The study should define the critical operations S_u from each use context before assessing coverage, document how

Table 30: Prospective responsibility-interface evaluation with comparison against an existing authorized audit.

Endpoint	Proposed independent evidence	Important rival explanation or harm
Forum identification	Mandates, referrals, reviewer reasons and case records	Competent route already exists but is underdocumented
Correction reach	Versioned decisions, transmission records, receiving-body disposition	Decision unchanged for sound evidential or legal reasons
Succession handling	Transfer instruments, record retention, access grants and residual powers	Programme change unrelated to representational use
Attribution discipline	Blind review of sampled contribution and authority claims	Increased false attribution or improper disclosure
Claimant burden	Safeguarded reports of access costs, delays, withdrawal and safety concerns	Self-selection and unobserved chilling effects

S_u may be contested, and record any operations added following protected, independently corroborated challenges. Treating the initially enumerated roles as the true universe of affected operations would risk reproducing the very exclusions under review.

For each sampled claim episode, independent reviewers could assess five separate outcomes: (i) a supported record of who carried out the material operation, (ii) a supported account of the legal or institutional competence governing that operation, (iii) the availability of a competent actor–forum answerability route, (iv) the capacity of that route to transmit an admissible correction or reasoned refusal to a consequential decision, and (v) a separately evaluated explanation of causal contribution and normative response, where such conclusions are appropriately within scope. Disagreements about any of these outcomes should be coded rather than converted to a binary judgment. An error in assigning a named individual responsibility must be treated as an adverse outcome, even if it improves nominal completeness of audit fields.

Table 30 summarizes the independent evidence and alternative explanations relevant to this prospective responsibility study.

The design should record non-registration, refusal of participation, and access restrictions as distinct sources of uncertainty. A protected reviewer might corroborate a missing forum while being legally unable to disclose the underlying testimony to the research team. This may yield a bounded procedural finding without establishing individual fault, revealing the confidential record, or permitting a claim-specific remedy. Evaluators should pre-specify a means of handling such partial findings, rather than making unrestricted disclosure an implicit condition of an audit’s success. Confidentiality breaches, retaliatory contact, adverse misattribution, and excessive demands on claimants belong among independently monitored harms.

Cases involving institutional succession require separate sampling or explicit stratification. Apparent gaps may be more common where custody, decisional power and review competence are dispersed; yet an ordinary records-transfer failure, changed retention rules, or a legally final decision could also explain the result. A useful comparison should ask whether the added interface improves the distinction among these conditions, rather than only whether it increases the number of actors named in a case file. The evaluation should allow for the possibility that less documentation, a protected referral and a clear public explanation offer a better justified outcome than exhaustive contribution tracking.

Negative findings must be reported. The intervention may duplicate an existing organizational chart, an established ombuds process, or an effective inter-agency review procedure. It may fail to provide access to any forum with competence to change the consequential use. It may increase costs and the exposure of personal information without improving independently corroborated dispositions. It may also classify a case as institutionally “covered” while its initial list of operations excludes material claimant experiences. Each would limit or defeat claims of incremental utility. As throughout this appendix, the endpoint is the comparative performance of an investigable procedure, not proof that an identified route provides substantive justice.

B.7 Evaluation of authority separation and remedy reach

A supplemental prospective evaluation module should test whether the authority-separated audit in Section 8.9 provides material information or improves independently justified dispositions beyond an effective established audit. Its central object is a *declared claim–response pair*: an independently eligible concern, its stated evidential question, a specified response type, and the institutions whose competence would be necessary to consider it. Each pair should be assigned a documented context and observation horizon; an asserted claim is not presumed legally valid merely because it is included in the research frame. Sampling should include category-scope objections alongside individual coding errors and should register the reasons cases are excluded. A review of only cases that successfully entered a formal appeal would be insufficient to establish broad accessibility.

The first measurement separates *claim reception* from *scope admission*. Investigators would ascertain whether a claimant had notice of an appropriate channel, could approach it without disproportionate procedural or safety costs, and could have a question about the governing representation accepted for consideration. The second separates *evidential reach* from *institutional competence*. Qualified reviewers would check, under lawful access procedures, which evidence was available, which operations the receiving forum could assess, and which independently documented mandates could connect its findings to a potential responding body. A third measurement concerns *action capability*: whether a specified legal, procedural, budgetary, or organizational power existed to consider a correction, prospective rule revision, reconsideration, explanation, or other stated action. The research must avoid inferring one capability from the presence of another.

For each claim–response pair, a structured record can therefore preserve at least four determinations: (i) reliable evidence established, unresolved, or unavailable; (ii) question within, outside, or ambiguously related to the reviewing forum’s scope; (iii) documented potential response route available, unavailable, or indeterminate; and (iv) actual response and independently evaluated consequence. Reasons for missing routes should be classified separately, including the absence of authority, lawful procedural finality, information access restrictions, insufficient resources, a protected third-party interest, and incomplete documentary evidence. No automated aggregation should rank these causes morally. Assessors should log uncertainty and contested interpretations, maintain role-specific access, and provide a protected disagreement channel for the people represented.

A comparator study could examine independently adjudicated differences in the recognition of material scope objections, accuracy of competence classifications, correct referrals to appropriately authorized fora, proportion of eligible claims obtaining a reasoned disposition, time and burden until resolution, unwanted disclosures, and substantive effects where these are independently measurable. An increased count of scope challenges could reflect improved observation without showing that the underlying institution became less just. Similarly, a referral recorded as delivered is not equivalent to an implemented remedy. The analysis should report endpoints by programme pathway and relevant claimant access conditions, while explicitly addressing selection into complaint submission and the possibility that protected reports cannot be included in an unrestricted dataset.

The comparative protocol should include *negative-control-like cases* whose review and remedy route is already clear, and cases with a legitimate restriction on retrospective change. The former test redundant documentation and possible burdens; the latter test whether the audit can distinguish a justified limit from missing institutional competence. There should also be a scenario in which a reviewer can obtain protected evidence but no public or cross-agency transfer is lawful. Correct classification of this case may require acknowledging uncertainty and preserving confidentiality rather than obtaining a larger number of record transfers. Any inference that a missing route *caused* an adverse substantive outcome requires a separately designed causal study; an authority map does not identify a counterfactual effect.

Before implementation, the research plan should determine how institutional mandates will be verified independently of the organization being evaluated, how conflicting legal interpretations will be treated, what opportunities exist to challenge the study’s own category choices, and who will adjudicate alleged researcher error or harmful disclosure. Protected community consultation may help reveal question types excluded by the institutional register. It may also impose costs and expectations that

the research is unable to satisfy. Safeguards should distinguish contribution to research from access to a remedy, and prohibit implying that participation itself creates legal standing or an entitlement to decision reversal.

The extension would be weakened if it merely documents a previously known separation, improves apparent route coverage by narrowing the set of admissible questions, or increases successful referrals without changing the quality or competence of decisions. It would be challenged more seriously by false legal-attribution claims, harm to vulnerable people, or substantial costs without corroborated benefit. These failure criteria should be predeclared. They preserve the distinction between a proposed diagnostic instrument and justified institutional reform, whose distribution of powers and obligations requires independent deliberation.

B.8 Prospective evaluation of contestability continuity

This module extends the authority-separation protocol with a study of documented claim transfer, reconstitution and competent review across institutional changes. The empirical object is a *claim-question-response episode*, which can cross representational versions or organizational boundaries. Eligible episodes would include independently identified challenges to category scope, evidential admission or consequential reliance, as well as routine corrections used for comparison. Configuration changes must be verified from institutional instruments and operational records. The proposed graph and handoff relations should be coded as a research representation of independently established arrangements; unverified authority is classified as unresolved. The method makes no presumption of a general entitlement to cross-boundary review.

Sampling should avoid conditioning exclusively on registered successful referrals. Where lawful and ethically approved, independent observation can include public guidance, protected interviews with intermediaries, oversight records, anonymized attempted referrals, and cases referred back as outside scope. Investigators should compare periods or institutional situations with credible baseline support, explicitly acknowledge selection into complaint channels, and predeclare which organizational changes define an eligible handoff episode. Consent and data-protection requirements constrain every proposed source. Constructed demonstration cases should remain separate from real observations.

For each relevant configuration, coders would extract the version-specific question and its declared permissible interpretations; the authorized claimant or intermediary route; any permitted record locator; and the competence to examine or act upon the requested issue. For each transfer, two independently assessed codings should distinguish *technical or documentary carriage* from *lawful and compatible claim continuation*. A coding of “route absent” requires adequate investigation of independently available substitute fora. “Undocumented”, “unverified”, “inapplicable” and “legitimately closed” should be retained as distinct statuses. Reviewers should record disagreements and make only the minimum protected evidence available to each role.

Proposed endpoints include (i) independently confirmed accuracy of route and jurisdiction classification, (ii) proportion of materially eligible claims with a safely usable route to a qualified reviewer, (iii) evidence that a compatible intermediate state exists across a multi-stage transfer, (iv) availability of a competent body to consider each declared response, (v) time, financial, learning and psychological costs of making a challenge, (vi) actual dispositions and independently judged reasons, and (vii) privacy incidents or suppression of participation associated with added linkage. These measures require independent denominators and confidence intervals when empirical estimation is feasible. Favorable outcomes and adverse consequences should be reported separately; an increase in complaint counts or favorable dispositions alone does not establish a just process.

A strong comparator would use an adequately functioning pre-existing appeal, archival or professional process with similar resources. Analyses could test whether the proposed handoff dossier contributes incremental detection of material claim losses or more accurate competent referrals once institutional mandate, case complexity, access to legal assistance, programme scale, and record quality are accounted for. A preliminary paper-based simulated assessment of documented cases might precede

ethically approved field research. Possible adverse results include higher exposure risk, mistaken attribution of jurisdiction, unnecessary re-identification, delay to other claimants, and diversion of staff from substantive service provision.

A prospective study should be interpreted conservatively if no safe external ascertainment is possible, if baseline bodies already provide equivalent competent review, or if applicable law forbids the linkage necessary for the proposed audit. In those settings, the continuity notation may clarify an interpretive question without constituting an implementable intervention. It cannot itself justify retaining personal information, override finality, create an appellate right, or supply normative agreement about an outcome.

B.9 Prospective assessment of competing challenges and remedial priorities

This module examines whether explicitly separating materiality, procedural urgency, institutional feasibility and remedial justification improves the quality and safety of consequential review. It remains an unimplemented study component, without recruited participants, observations, or estimated effects. The intervention is a structured, question-indexed deliberation record supplementing an established authorized workflow. The comparator is a mature procedure for triage, review, rights protection and reasoned disposition, including existing priority-setting practices where applicable. An evaluated institution must first document the legal and ethical basis on which any allocation or remedial decision may be made. The added audit is not an independent source of jurisdiction.

Eligibility and unit of comparison. The primary unit for this module is a *claim–candidate–response deliberation episode* within a documented institutional configuration. Episodes should be nested within institutions, pathways and capacity settings for analysis, with safeguards against double-counting one underlying complaint as several independent observations. Sampling should include independently ascertained eligible claims outside routine registries when lawful, proportionate and safe; a system restricted to accepted complaints cannot support conclusions about access for persons whose claims could not be filed. An institution may make several decisions concerning one question, including a procedural scheduling decision and a distinct substantive disposition. Each should retain its own authority, evidential horizon, grounds and effects.

Predeclared observations and adjudication. For each included episode, qualified independent reviewers should record: the asserted material harm and its epistemic status; the reason for any urgent scheduling; available and excluded response options; the legal and institutional warrant for each feasibility conclusion; possible third-party interests; evidence-access restrictions; the declared evaluation dimensions and limits on their comparability; uncertainty about plausible continuation scenarios; the actual review or remedial disposition; adverse consequences; and the options for appeal or revision. Missing, inapplicable, unverified and legally protected fields should be coded separately. Evaluation should record whether a decision was explained, *not* classify an action as unjust merely because it was refused. The existence of an adverse outcome should not by itself prove that the intervention caused it.

Table 31 identifies the proposed measures and negative cases for the remedial-priority module.

Analysis and comparison. The evaluation should prespecify how institutional clustering, different triage mandates, changing capacity, claimant selection and safety restrictions will be addressed. A proposed primary supplementary endpoint is the proportion of independently eligible episodes receiving a timely, reasoned disposition through a verified competent route *without an independently adjudicated serious rights incident*, with numerator, denominator, horizon and event definitions fixed before data collection. It should be reported alongside urgent-protection access, unmet eligible needs, qualitative accounts of contested values, time costs, disclosure incidents, and effects on other eligible claimants. These quantities have different denominators and interpretations. An aggregate figure may be useful for comparison but should not be presented as a metric of overall ethical justice. Where designs permit causal

Table 31: Proposed measures and falsifying observations for the remedial-priority module.

Domain	Candidate measure or review	Important limitation or negative case
Triage and protection	Time or path to an authorized provisional safeguard, conditional on verified eligibility	Slower protection despite improved documentation
Reason-giving	Independent assessment of the relevance, intelligibility and reviewability of stated grounds	Reasons already sufficient in standard practice
Institutional authority	Verified response competence and recorded prohibitions or uncertainty	Apparent favorable outcomes achieved outside lawful powers
Plural consequences	Disaggregated outcomes for claimants, others and service capacity	Aggregate gain disguises displacement or rights injury
Safe participation	Off-register ascertainment, burdens, protected routes and assistance	Easier access only for well-documented claimants
Uncertain comparisons	Frequency and handling of incomparable or unresolved options	False certainty or indecision introduced by the audit
Reconsideration	Independent availability of reasoned review of the procedure and substantive choice	Appeals delayed or obstructed by additional paperwork

estimation, any causal claim requires explicit intervention assignment or identification assumptions and sensitivity analysis; otherwise results remain descriptive.

Contrary findings, stopping criteria and boundaries. A result against the intervention may occur when established reviewers already recognize the necessary distinctions, when augmented recording defers urgent action, when supposed robustness suppresses lawful discretion, when an ostensibly non-dominated option causes an overlooked third-party rights violation, or when participants encounter new learning, psychological or compliance burdens. The study should define independent safety review and procedures for suspending an intervention whose recordkeeping causes serious harm. Records concerning protected persons must not be made public as a price of reproducibility. External validation requires qualified ethical and legal review, affected-party engagement where appropriate, and documentation of reasonable dissent. An appropriately authorized and reasoned decision may remain contestable; that fact is not evidence of procedural failure by itself.

B.10 Prospective evaluation of evaluative-frame governance

This unimplemented module asks whether a structured audit of changes in question scope, participant standing and evaluative criteria improves the competent treatment of materially relevant challenges relative to established institutional rule-review procedures. The unit of analysis is a *frame-challenge episode*, defined by a documented request to retain, amend or independently review a specified question, admission rule, evaluative dimension, or division of institutional authority. Episodes concerning mere correction of an entry should remain distinguishable, even if they motivate a broader scope inquiry. Inclusion cannot depend exclusively on the body whose criterion is being challenged. Instead, a lawful independent sampling channel would attempt to ascertain off-register attempts, protected referrals and rejected scope complaints, with an explicit account of persons whose participation cannot safely be assessed. Such procedures require context-specific consent, ethics and legal authorization; they are proposals, not operational methods already tested.

Independent ascertainment and protected denominators. For an authorized sampling frame S , a primary descriptive comparison may consider the proportion of independently eligible frame-challenge episodes receiving a competent, reasoned disposition within a predeclared horizon. The denominator requires independent adjudication of eligibility and may not be estimated merely from the administrative register. Reports should give the separately established count of eligible observed episodes, a defensible account of unobserved cases where available, the fraction whose eligibility could not be determined,

Table 32: Unimplemented assessment module for challenges to evaluative frames.

Dimension	Prospective observation	Negative, uncertain or adverse result
Safe ascertainment	Protected independent capture of registered and off-register scope challenges	No lawful or defensible off-register denominator
Challenge competence	Verified rule-review authority and route for disputes about standing	Referral without actual power to consider the question
Epistemic engagement	Reasons received, permissible disclosure, independent assessment and documented dissent	Formal participation without engagement or unsafe demands for testimony
Rule change	Explicit old and revised versions, type of revision, treatment of pending claims	Ambiguous migration, covert retrospective use, or lawful closure misclassified as failure
Distributional effect	Outcomes for newly recognized, previously admitted and indirectly affected persons	Benefits for documented parties accompanied by displacement or stigma
Review burden	Learning, compliance and psychological costs, delays, confidentiality incidents	Increased exposure, intimidation or lower access despite better records
Comparator	Established authorized appeal, policy consultation or independent ombud process	Existing procedure performs at least as well with fewer costs

and the method by which safe exclusions were handled. It may be impossible to identify an overall participation rate without unjustified inference about unobserved persons. Disposition is distinct from acceptance of a challenge, subsequent amendment, and improvement in consequential outcomes. Each uses a separately declared interpretation and denominator.

Table 32 records candidate observations and limitations for evaluative-frame governance.

Intervention, comparators and identification. The proposed intervention is the addition of a typed frame-change register, a protected route for scope questions where independently authorized, and independent checking of rule-version transitions. The principal comparator is a mature, appropriately staffed system already able to review changes in institutional criteria, not the absence of review. If assignment is feasible and authorized, a clustered, staggered or matched comparison may be considered subject to safeguards against untreated urgent harm. No randomization or causal identification is presumed. Changes in legal powers, service budgets, reviewer expertise, outreach, enforcement priorities, advocacy and the composition of the claimant population are rival explanations. Contextual changes that materially revise the comparator or the admissible question should be documented as such. Qualitative process reconstruction should test whether the additional audit information actually reached a competent decision maker and affected reasons or dispositions.

Predeclared adverse cases and interpretation. The module must retain observations in which independent review finds no supported recognition gap, legitimate standing restrictions, confidentiality-based non-disclosure, or warranted refusal to amend an established criterion. It must also record whether frame revision creates a stigmatizing category, undermines a protected relationship, delays urgent action, or diverts scarce review capacity. Where an affected party declines safe participation, the study should not classify the absence of an interview as evidence that no interest exists. Conversely, speculative assertions of universal affectedness should not override rules of relevance or legitimate jurisdiction. Findings should be reported as documentary access, adjudicated competence, procedural decision, and observed consequence, rather than compressed into an overall score of democratic legitimacy or justice. This module reports no participant recruitment, empirical measurement, evaluated policy or legal determination.

B.11 Prospective study of recursive rule-review and meta-contestation

This unimplemented module examines whether the proposed differentiation of first-order complaints, challenges to standing, and challenges to the revision procedure yields demonstrably better institutional

Table 33: Planned observables and adverse cases for recursive rule-review.

Dimension	Proposed observation	Contrary or limiting evidence
Challenge type	Independent classification of data, scope, revision-rule and authority objections	Ambiguous levels or error-prone classifier
Reachability	Documented submission, lawful referral and capacity for substantive consideration	Nominal referral without competence or protected access
Reason quality	Independently assessed response to the challenged rule and material evidence	Formulaic response or selective engagement
Rule lineage	Versioned authorizations, effective scope, pending-case policy, provenance	Undocumented change or false causal attribution
Affected interests	Protected off-register ascertainment, third-party rights and practical costs	No safe denominator, displacement, disclosure or retaliation
Comparator	Existing appeal, consultation or rule-amendment process	Existing mechanism performs at least as well
Outcome	Disposition, rule change and subsequent consequences separately coded	Increased delay, unlawful action, negligible benefit

review than an established competent comparator. Its unit of analysis is a *meta-challenge episode*: a documented attempt, including a protected or off-register attempt where lawfully ascertainable, to challenge a rule governing which claims, evidence or revisions may be considered. A challenge may be sincere but unfounded, or well supported but outside the programme's lawful competence. The study must record those possibilities separately. A completed consultation and a competent consideration of a challenged admission rule are not interchangeable endpoints.

Candidate contrasts and ascertainment. An appropriately staffed existing statutory review or rule-making arrangement is the primary comparator. The proposed intervention is an optional typed record of the rule challenged, the applicable level of admission, the actual forum's authority, the evidential protections used, and the trace of reasons and effects across versions. Study eligibility and sampling channels would be defined in consultation with independent ethics and legal reviewers. Records of rejected portal submissions, lawfully accessible referral logs, and consent-based independent submissions might complement admitted cases where justified. Safe non-participation must not be interpreted as absence of affectedness; protected claimant information must not be reproduced merely to improve the denominator. If eligibility beyond the observed population cannot be independently assessed, no system-wide recognition rate should be reported.

Table 33 specifies proposed observations and contrary evidence for recursive rule-review.

Identification, independent judgment and stopping conditions. A study should predeclare the observation window relative to institutional decision sequences, analytic eligibility rules and the question against which each outcome is assessed. The concept of $\text{Reach}_{\text{int}}$ in Section 8.13 is a logical comparison set rather than an observed measure of governmental openness. Its empirical analogue would require verification of actual decision powers and permissible revision paths, which may be partly unavailable. Potential confounders include statutory reform, changes in leadership, civic mobilization, review staffing, judicial involvement, resource reallocation, and contemporaneous reporting changes. Because a rule change can alter which challenges are observable, comparisons must not treat the new administrative denominator as stable without independent warrant. Sensitivity analysis should distinguish non-detection from supported absence of a route.

The prospective assessment should include qualitative examination of negative and mixed cases: a competent existing reviewer already processes meta-challenges; a new channel exposes vulnerable claimants; an apparently successful rule change was produced by another authorized process; a proposed expanded standing rule conflicts with a protected third-party interest; or indiscriminate reopening disrupts an appropriately finalized proceeding. Stopping and escalation criteria should address retaliation, involuntary disclosure, jurisdictional uncertainty, and undue delays to urgent duties. No programmatic

Table 34: Illustrative integrated-register endpoints for a prospective comparison, with independent assessment and explicit negative interpretations.

Endpoint	Proposed verification	Alternative or adverse interpretation
Material omission identified	Independent adjudication against a defended question and protected comparison record	Changed ascertainment, contested question or invalid coding
Effective evidential access	Verified permission and secure receipt by a competent reviewer	Disclosure risk, unauthorized access or missing source evidence
Competent case disposition	Governing mandate, reasoned response and distinction between acknowledgement and reconsideration	Nominal review or unverified jurisdiction
Transmission to consequential use	Matched case and policy versions, documented receipt and implementable decision path	Corrected archive without effective decision reach
Rule-level challenge	Independently classified scope objection and competent authorized rule review	Closed admission frame or justified protective closure
Adverse effects and displacement	Protected incident assessment, third-party burdens and resource-use evidence	Retaliation, delay, confidentiality breach or diverted safeguards

audit creates an entitlement to legal review or a waiver of protective restrictions. Actual intervention, access and research recruitment would require separate approval; this appendix reports no participants, measured outcomes or causal estimates.

B.12 Integrated claim-pathway register and interpretation

The synthesis in Section 9.19 requires an assessment record that preserves differences among kinds of evidence and institutional competence. For each independently eligible claim episode, an evaluator should construct a versioned seven-field profile corresponding to Equation 28. Each field should record (a) its coding value, (b) the source and access status of its evidential warrant, (c) the responsible independent assessor, (d) disagreement or uncertainty, and (e) the relevant law, rule or institutional version. The profile is a planning and audit instrument, not an observed ethics score or a requirement for public exposure of protected material.

Table 34 sets out the proposed comparative endpoints for the integrated warrant register.

Comparisons should use an existing competent audit and appeal process as the principal comparator, with variation in institutions and claim types declared independently of observed outcomes. A register that contains more records can perform worse if it imposes unsafe disclosure, omits unregistered claimants, misstates competence, or consumes capacity needed for urgent protection. Likewise, a low transfer rate need not indicate misconduct when protected finality or a lawful limit on reopening is documented. Investigators should report field-specific estimates, unresolved evidence, disagreement rates and protected qualitative accounts. An aggregate “complete pathway” percentage may be supplied as descriptive metadata only when its selection rule and denominator are explicit; it must not stand in for an ethical outcome.

A formative study might examine coding agreement and feasibility before testing institutional effects. Any subsequent causal comparison requires an intervention contrast, eligible population, independently assessed outcomes, plausible counterfactual support and a declared follow-up scheme. Registration, legal review, ethical approval and data protection decisions remain prerequisites for a field implementation. The current appendix provides no observations, causal estimates or implemented interventions.

B.13 Analytical plan and rival explanations

The analysis should first present transparent case flow, institutional variation, source completeness, and differences between the standard and augmented workflows. A predeclared primary comparison should be reported with uncertainty intervals appropriate to the design and the sampling units, avoiding false precision from clustered cases or selective follow-up. Any causal estimate must explicitly state the population, the intervention contrast, the target outcome, and the assumptions required for identification. Variation across claimant groups and institutional pathways may be as important as an aggregate mean. Differences in materiality criteria or evidential access should be reported as changes to the interpretive setting, rather than silently absorbed into an overall success score.

A qualitative process-tracing component should reconstruct selected decisions from initial record formation through downstream review. It should seek evidence that the supplemental interface revealed a missing distinction and caused a competent institution to engage with that distinction, while considering rival pathways. Particularly relevant alternatives are strengthened enforcement, a new legal rule, increased funding, independent community advocacy, and personnel changes. Evidence that the same distinction was already identified and acted upon under the standard audit would reduce the claimed added value of the interface. Evidence of greater administrative burden or avoidable disclosure may count against its use even if additional omissions are detected.

The six OECD evaluation lenses of relevance, coherence, effectiveness, efficiency, impact, and sustainability can inform the selection of evaluative questions and the interpretation of costs and consequences [35]. The OECD criteria are contextual evaluation lenses, not a method that automatically supplies causal identification or an ethical decision rule. A study may therefore reasonably find the proposed audit informative but inefficient, effective for correction delivery but weak in long-term outcomes, or valuable to reviewers while burdensome to claimants. Such mixed findings belong in the research record.

B.14 Predeclared conditions limiting the contribution

The evaluation should treat the following outcomes as potential evidence against the proposed incremental contribution: a competent existing audit identifies the same material omissions with comparable evidential quality; its corrections reach the same authorized reviewers with equivalent timeliness; the augmented procedure systematically misclassifies contextual claims; apparent gains disappear after accounting for better staffing, case selection, or expanded legal powers; the workflow causes harmful disclosure or retaliation; or its costs displace other protective and remedial activities. The first two observations would suggest redundancy, the next two challenge validity and attribution, and the final two raise independent ethical and institutional objections.

A further limiting case arises where relevant lived experiences cannot be represented by the declared classifications without material distortion. Additional fields or more detailed records may worsen exposure or reinforce an institutionally imposed vocabulary. The appropriate response may be protected qualitative inquiry, participation in the choice of categories, or suspension of a consequential model, subject to competent authorization. This possibility limits claims that a refined quantitative record can resolve representational injustice by itself. It is consistent with established warnings about the loss of sociotechnical context in abstraction [43].

Results should be retained and reported when they reveal equivalence, harm, inconclusive causal attribution, or a failure to identify the proposed mechanism. Publication of a negative result is part of evaluating the framework, not a deviation from its research objective.

B.15 Governance, protection, and reproducibility

Prior to any field implementation, the investigators would need approval from the relevant institutional and ethics bodies, lawful grounds for access and processing, a data-management plan, and agreements establishing reviewer independence and secure handling of protected records. Risks include reputational

injury, identification of vulnerable claimants, exposure of contested testimony, coercion through institutional dependence, and retaliation against participants. A research protocol should specify consent and withdrawal where applicable, retention and access limits, independent complaints routes, the handling of legal confidentiality, and criteria for suspending data collection or consequential deployment. The study must avoid treating public disclosure as the default remedy for inadequate transparency.

Reproducibility should be sought through versioned protocols, de-identified metadata where lawful, a transparent coding dictionary, archived analytic scripts, an independent record of modifications, and separate reporting of deviations from the predeclared analysis. Case-level raw evidence and identities need not be made public for a study to be scrutinizable, provided secure independent audit and lawful procedures are available. Where even protected verification is impossible, the evidential limitation should be stated rather than replaced by an assertion of completeness. PROV-O may assist with representing the lineage of documentary assertions; its provenance vocabulary does not certify their truth or ethical acceptability [47].

The accompanying `protocol/study_specification.json` is a machine-readable *planning template*; it contains no observed case data and makes no registration claim. Its declared outcomes and methodological warnings may support later preregistration, local adaptation, and independent review. The schema should be frozen only after context-specific study design, legal analysis, ethical approval, and author review. Actual study observations and findings would require a separate future research report.

B.16 Interpretive boundary

The proposed evaluation investigates whether a cross-stage accountability interface improves the documentation and competent review of material claims relative to a meaningful comparator. It does not substitute model adequacy for moral legitimacy, receipt of notice for effective remedy, or detected omission for demonstrated causal harm. The protocol’s contribution is to state in advance what evidence would support, weaken, or leave indeterminate its practical claim. Empirical fieldwork, independent assessment, and normative argument remain separate requirements for deciding whether the interface deserves adoption in a particular institutional setting.

C Contextual Transfer Records for the Four-Paper Inquiry

This appendix provides an audit instrument for claims that combine the four manuscripts’ distinct research objects. Its purpose is to make the boundaries of an inference inspectable at the moment of transfer. The companion series supplies conceptual questions and proposed analytical devices; it supplies no independent confirmation of the empirical or normative conclusions examined here. The record introduced below is an editorial and methodological device rather than a technical standard, an enforceable right, or a novel mathematical result.

C.1 The unit of transferred inference

A transfer begins with an identifiable statement, rather than a shared keyword. The relevant unit is a *claim episode*: an assertion concerning a specified historical domain, a question addressed to that domain, the representation through which the assertion was obtained, and the grounds invoked to interpret it. Equation 29 records a proposed audit tuple. Its components are descriptive placeholders whose meaning is established by the local inquiry.

$$T = (D, Q, S, E, A, P, O, W). \quad (29)$$

Here D is the domain of admitted or independently reconstructed histories; Q is the declared interpretive question; S is the representation or sequence of reduction maps; E is the evidential record and its

Table 35: Working concordance for transferring claims across the four research objects.

Expression	Working role in the series	Unwarranted substitution to test
Generative capacity	Historically constituted possibilities for continued contribution, relation, and development, examined especially in Paper I.	A high recorded output or stable aggregate does not establish maintained capacity.
Credible commitment	An institutional promise whose reliance and review conditions are investigated in Paper II.	A corrected data field, formal pledge, or registered complaint does not alone establish credible implementation.
Normative claim	An interpretable question about obligations, standing, or consequences over admitted histories, investigated in Paper III.	A correctly answered declared query does not settle whether the query family was ethically adequate.
Representational loss	A query-relevant distinction absent from a specified reduction, analyzed in Paper IV.	Loss within a stipulated model does not alone establish an observed harm, unlawful act, or moral wrong.
Remedy	A legally or ethically grounded response through an applicable institutional path.	Review, reversal, correction, and improved substantive conditions are separate possible outcomes.
History and scale	Declared structures of events, relations, interpretation, and reduction, varying with the research task.	Record timestamps and aggregate administrative levels do not establish a universal time order or a unique normative hierarchy.

provenance; A is an account of institutional competence and access; P specifies the available procedural path, where one exists; O is the observed or stipulated outcome; and W states the warrant type and remaining uncertainties. These fields are not naturally complete or universally enumerable. They make the assumptions of a particular transfer explicit and open to revision.

A transfer from Paper III's representational result to Paper IV's account of institutional responsibility, for example, needs the reduction S , the admissible histories D , and the query Q to interpret an information-loss witness. It additionally needs evidence in E about the construction and use of the representation. A further transition to Paper II's question of an effective remedy requires an independently supported account of competence A , the procedural path P , and its observed disposition O . A claim that the correction changed generative capacity in the sense investigated by Paper I calls for a historical and causal assessment beyond these documentary records. The tuple provides a record of the successive inquiries; it does not prove that one of them entails the next.

C.2 Terminological concordance

Table 35 records recurring expressions at the level of their principal research roles. Its distinctions are deliberately narrower than a complete theory of value, obligation, or social ontology. In particular, institutional records may be used to investigate relations while remaining imperfect and historically contingent access points to those relations.

The term *scale* is especially susceptible to drift. A finer documentary classification, a more inclusive claimant population, a more detailed institutional process, and a physical change of resolution are different operations. A representation may gain detail while still omitting an excluded history. A change in the admissible population is a domain change, whereas a change in query interpretation can occur without changing the record. A physical renormalization flow entails additional mathematical structures absent from the ordinary institutional comparisons in this study. Claims about these operations should identify which transformation has occurred.

C.3 Two constructed transfer audits

A verified classification correction with unchanged allocation. A hypothetical office corrects an erroneous eligibility category in a record after a participant supplies independent documentation. The office retains the original category, the evidence supporting revision, and the corrected version. A separate allocation committee receives the revised record, considers it under an authorized procedure, and gives reasons for maintaining the original allocation under a criterion independent of the corrected category. The record supports a classification error, a correction, actual transmission, and a reasoned review. It does not support a claim that allocation changed, that an effective substantive remedy was granted, or that generative capacity was restored. Further ethical inquiry may criticize the allocation rule, but that criticism needs premises and evidence beyond the correction record. The institutional path therefore counts as operational review while the outcome remains unchanged.

An accurate aggregate with an off-register ecological claim. A second hypothetical report accurately summarizes admitted production records. An affected community provides independently collected evidence of an ecological burden that the report's admission rules exclude. The aggregate can be correct relative to its input domain while remaining uninformative about the ecological question. The responsible inquiry first examines whether the claimed burden falls within the report's authorized purpose, whether the exclusion affected a consequential use, and what evidence an appropriate reviewer may lawfully receive. If a competing environmental review already covers the concern and permits meaningful participation, an additional cross-stage audit may offer little marginal benefit. If neither review can receive the evidence, the proposed framework identifies a possible gap in representation and procedure; the existence and normative significance of an ecological harm remain matters for separate assessment.

Both cases are constructed. The audit distinguishes negative outcomes for a proposed causal mechanism from failures of data collection and from limits imposed by valid confidentiality rules. An observed difference in an administrative record, without evidence of its production and practical use, is not sufficient to classify the entire chain.

C.4 Documentary audit and comparative limits

A transfer record should identify the version of the source statement, its local section or proposition, the admissible domain and question, the form of evidence actually consulted, and the reason an additional warrant is needed. The companion archive supplies a machine-readable cross-paper concordance with bounded manuscript-level locations and a distinction between internal text inspection and external source verification. The concordance records research objects and candidate misreadings; it is not a claim that every passage in the series has been independently verified.

For documentary reconstruction, the W3C provenance vocabulary provides established means of describing entities, activities, and derivations [47, 46]. Provenance statements facilitate tracing relevant assertions but do not certify the truth of a contested event. The OECD evaluation criteria similarly offer context-sensitive perspectives on an intervention; they do not independently identify the causal effect of this manuscript's protocol [35]. Applicable legal duties must be traced to the relevant law and competent authority, with the safeguards for automated decisions under the GDPR remaining confined to their legal scope [14]. None of these sources supplies a universal institutional inference rule.

A useful test of the concordance is therefore comparative. A competent conventional institutional audit may already recover the relevant distinctions, establish the same limitations, and furnish the same review path. In such a case the present transfer record may improve traceability but offer negligible independent explanatory value. Conversely, a documented case in which the record reveals a previously overlooked admissibility or propagation failure could support a bounded claim of additional diagnostic usefulness. The prospective study in Appendix B would be required to move from such illustrations to comparative evidence. Its implementation and outcomes remain open.

Research Ethics and Source Provenance

This manuscript is an exploratory conceptual and formal preprint candidate prepared for author review. Appendix A contains a reproducible constructed case that establishes internal logical consistency only. Appendix B sets out an unimplemented prospective comparison protocol; Appendix C provides a terminology and inferential-transfer audit for the four-paper series; no field study or empirical result is reported. Cases are stipulated constructions; they describe no identifiable subjects or actual institutional results. The elementary fiberwise criterion is inherited from established set theory and the author’s prior exploratory two-complex work. Mathematical adequacy is assessed only relative to declared histories, interpretations, and questions. Historical and institutional truth, causal adequacy, legitimacy, and moral justification require further independent argument and evidence. The bibliographic verification ledger distinguishes publisher abstracts, standards documents, institutional reports, and open-access scholarly text. The illustrative institutional comparisons in Sections 2–9 are constructed rather than empirical findings. The fourteen conditional propositions and integrated warrant register establish neither universal standing, legal review powers, a complete roster of affected persons, nor ethical legitimacy through procedural consistency. The manuscript requires further passage-level source verification, author approval, and independent scholarly assessment before dissemination.

OpenAI’s ChatGPT assisted with source discovery, comparative analysis, English drafting, LaTeX authoring, and build checks. The human author is responsible for the manuscript’s accuracy, approval, verification, and dissemination. Sources cited in this draft imply no endorsement. Copyright © 2026 Wanhong Huang. The author proposes to distribute the original textual material under Creative Commons Attribution–NonCommercial 4.0 International (CC BY-NC 4.0) upon approval, subject to third-party rights.

References

- [1] Mike Ananny and Kate Crawford. “Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability”. In: *New Media & Society* 20.3 (2018), pp. 973–989. DOI: 10.1177/1461444816676645.
- [2] Elizabeth Anderson. *Value in Ethics and Economics*. Harvard University Press, 1993. ISBN: 9780674931893.
- [3] Elizabeth Anderson. “The Epistemology of Democracy”. In: *Episteme* 3.1–2 (2006). Epistemic diversity and feedback in democratic learning; institutional analogies are qualified, pp. 8–22. DOI: 10.3366/epi.2006.3.1–2.8.
- [4] Sherry R. Arnstein. “A Ladder of Citizen Participation”. In: *Journal of the American Institute of Planners* 35.4 (1969), pp. 216–224. DOI: 10.1080/01944366908977225.
- [5] David Blackwell. “Equivalent Comparisons of Experiments”. In: *The Annals of Mathematical Statistics* 24.2 (1953), pp. 265–272. DOI: 10.1214/aoms/1177729032.
- [6] Mark Bovens. “Analysing and Assessing Accountability: A Conceptual Framework”. In: *European Law Journal* 13.4 (2007), pp. 447–468. DOI: 10.1111/j.1468-0386.2007.00378.x.
- [7] Geoffrey C. Bowker and Susan Leigh Star. *Sorting Things Out: Classification and Its Consequences*. MIT Press, 1999. DOI: 10.7551/mitpress/6352.001.0001.
- [8] Donald T. Campbell. “Assessing the Impact of Planned Social Change”. In: *Journal of MultiDisciplinary Evaluation* 7.15 (2011). Reprint of Occasional Paper Series no. 8 (1976), pp. 3–43. DOI: 10.56645/jmde.v7i15.297.
- [9] Catherine D’Ignazio and Lauren F. Klein. *Data Feminism*. MIT Press, 2020. ISBN: 9780262044004. URL: <https://data-feminism.mitpress.mit.edu/>.
- [10] Norman Daniels and James E. Sabin. “Accountability for Reasonableness”. In: *Setting Limits Fairly: Can We Learn to Share Medical Resources?* The four procedural conditions arise in the institutional context of health-care priority setting. Oxford University Press, 2002, pp. 43–66. DOI: 10.1093/acprof:oso/9780195149364.003.0004.
- [11] Kristie Dotson. “Tracking Epistemic Violence, Tracking Practices of Silencing”. In: *Hypatia* 26.2 (2011), pp. 236–257. DOI: 10.1111/j.1527-2001.2011.01177.x.
- [12] Ron Eglash et al. “Computational Reparations as Generative Justice: Decolonial Transitions to Unalienated Circular Value Flow”. In: *Big Data & Society* (2024). DOI: 10.1177/20539517231221732.
- [13] Wendy Nelson Espeland and Michael Sauder. “Rankings and Reactivity: How Public Measures Recreate Social Worlds”. In: *American Journal of Sociology* 113.1 (2007), pp. 1–40. DOI: 10.1086/517897.
- [14] European Parliament and Council of the European Union. *Regulation (EU) 2016/679 (General Data Protection Regulation)*. Article 22, with context-specific statutory safeguards. 2016. URL: <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>.

- [15] Nancy Fraser. “Abnormal Justice”. In: *Critical Inquiry* 34.3 (2008). Used specifically for contestation of the frame of justice, including questions of standing, pp. 393–422. DOI: 10.1086/589478.
- [16] Miranda Fricker. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press, 2007. DOI: 10.1093/acprof:oso/9780198237907.001.0001.
- [17] Archon Fung. “Varieties of Participation in Complex Governance”. In: *Public Administration Review* 66.s1 (2006), pp. 66–75. DOI: 10.1111/j.1540-6210.2006.00667.x.
- [18] Timnit Gebru et al. “Datasheets for Datasets”. In: *Communications of the ACM* 64.12 (2021), pp. 86–92. DOI: 10.1145/3458723.
- [19] Pamela Herd and Donald P. Moynihan. *Administrative Burden: Policymaking by Other Means*. Russell Sage Foundation, 2018. ISBN: 9780871544445. DOI: 10.7758/9781610448789.
- [20] Miguel A. Hernán and James M. Robins. *Causal Inference: What If*. Author-maintained online book; consulted for the specification of causal questions and target-trial emulation. Boca Raton: Chapman & Hall/CRC, 2020. URL: <https://miguelhernan.org/whatifbook>.
- [21] International Organization for Standardization and International Electrotechnical Commission. *ISO/IEC 42001:2023 Information Technology—Artificial Intelligence—Management System*. Tech. rep. ISO/IEC 42001:2023. ISO/IEC, 2023. URL: <https://webstore.iec.ch/en/publication/90574>.
- [22] IPCC. *Climate Change 2022: Impacts, Adaptation and Vulnerability. Summary for Policymakers*. Intergovernmental Panel on Climate Change, 2022. URL: <https://www.ipcc.ch/report/ar6/wg2/chapter/summary-for-policymakers/>.
- [23] Sheila Jasanoff. “Technologies of Humility: Citizen Participation in Governing Science”. In: *Minerva* 41.3 (2003), pp. 223–244. DOI: 10.1023/A:1025557512320.
- [24] Margot E. Kaminski and Jennifer M. Urban. “The Right to Contest AI”. In: *Columbia Law Review* 121.7 (2021), pp. 1957–2048. URL: <https://www.columbialawreview.org/content/the-right-to-contest-ai/>.
- [25] Carl Knight. “Reflective Equilibrium”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Methodological comparison; does not certify participatory institutional legitimacy. Metaphysics Research Lab, Stanford University, 2023. URL: <https://plato.stanford.edu/entries/reflective-equilibrium/>.
- [26] Joshua A. Kroll et al. “Accountable Algorithms”. In: *University of Pennsylvania Law Review* 165.3 (2017). URL: <https://pennlawreview.com/2017/02/23/accountable-algorithms/>.
- [27] Leslie Lamport. “Time, Clocks, and the Ordering of Events in a Distributed System”. In: *Communications of the ACM* 21.7 (1978), pp. 558–565. DOI: 10.1145/359545.359563.
- [28] Hans Lindahl. *Fault Lines of Globalization: Legal Order and the Politics of A-Legality*. Boundary and normative-responsiveness discussion limited to publisher available book and chapter abstracts. Oxford University Press, 2013. DOI: 10.1093/acprof:oso/9780199601684.001.0001.
- [29] Christian List and Laura Valentini. “What Normative Facts Should Political Theory Be About? Philosophy of Science Meets Political Liberalism”. In: *Oxford Studies in Political Philosophy*. Vol. 6. Oxford University Press, 2020, pp. 185–220. DOI: 10.1093/oso/9780198852636.003.0007.
- [30] José Medina. *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and the Social Imagination*. Oxford University Press, 2013. DOI: 10.1093/acprof:oso/9780199929023.001.0001.
- [31] National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Tech. rep. NIST AI 100-1. National Institute of Standards and Technology, 2023. DOI: 10.6028/NIST.AI.100-1.
- [32] Helen Nissenbaum. “Accountability in a Computerized Society”. In: *Science and Engineering Ethics* 2.1 (1996), pp. 25–42. DOI: 10.1007/BF02639315.
- [33] Helen Nissenbaum. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2010. ISBN: 9780804752374.
- [34] OECD. *OECD Due Diligence Guidance for Responsible Business Conduct*. Paris: OECD Publishing, 2018. DOI: 10.1787/15f5f4b3-en.
- [35] OECD. *Better Criteria for Better Evaluation: Revised Evaluation Criteria Definitions and Principles for Use*. Paris: OECD Publishing, 2019. DOI: 10.1787/15a9c26b-en.
- [36] OECD. *Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449)*. Adopted 22 May 2019; amended 3 May 2024; non-binding recommendation; Section 1.3(iv). 2024. URL: <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>.
- [37] Elinor Ostrom. “Polycentric Systems for Coping with Collective Action and Global Environmental Change”. In: *Global Environmental Change* 20.4 (2010). Polycentric climate-governance comparison, not a general warrant for institutional legitimacy, pp. 550–557. DOI: 10.1016/j.gloenvcha.2010.07.004.
- [38] Philip Pettit. “Republican Forms: Constitutionalism and Democracy”. In: *Republicanism: A Theory of Freedom and Government*. Publisher chapter abstract consulted; constitutional account not equated with a generalized administrative mandate. Oxford University Press, 1999, pp. 171–205. DOI: 10.1093/0198296428.003.0007.
- [39] Paul Pierson. “Increasing Returns, Path Dependence, and the Study of Politics”. In: *American Political Science Review* 94.2 (2000), pp. 251–267. DOI: 10.2307/2586011.
- [40] Inioluwa Deborah Raji et al. “Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (2020), pp. 33–44. DOI: 10.1145/3351095.3372873.

- [41] Charles F. Sabel and Jonathan Zeitlin. “Experimentalist Governance”. In: *The Oxford Handbook of Governance*. Ed. by David Levi-Faur. Publisher abstract supports iterative governance in specified contexts only. Oxford University Press, 2012, pp. 169–184. DOI: 10.1093/oxfordhb/9780199560530.013.0012.
- [42] Andrea Saltelli et al. “Five Ways to Ensure That Models Serve Society: A Manifesto”. In: *Nature* (2020). DOI: 10.1038/d41586-020-01812-9.
- [43] Andrew D. Selbst et al. “Fairness and Abstraction in Sociotechnical Systems”. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency* (2019), pp. 59–68. DOI: 10.1145/3287560.3287598.
- [44] Cass R. Sunstein. “Incommensurability and Valuation in Law”. In: *Michigan Law Review* 92.4 (1994), pp. 779–861. URL: <https://repository.law.umich.edu/mlr/vol92/iss4/2/>.
- [45] Dennis F. Thompson. “Moral Responsibility of Public Officials: The Problem of Many Hands”. In: *American Political Science Review* 74.4 (1980), pp. 905–916. DOI: 10.2307/1954312.
- [46] W3C Provenance Working Group. *PROV-DM: The PROV Data Model*. W3C Recommendation. World Wide Web Consortium, 2013. URL: <https://www.w3.org/TR/prov-dm/>.
- [47] W3C Provenance Working Group. *PROV-O: The PROV Ontology*. W3C Recommendation. World Wide Web Consortium, 2013. URL: <https://www.w3.org/TR/prov-o/>.
- [48] Kenneth G. Wilson and John Kogut. “The Renormalization Group and the ϵ Expansion”. In: *Physics Reports* 12.2 (1974), pp. 75–199. DOI: 10.1016/0370-1573(74)90023-4.
- [49] World Bank. *World Development Report 2021: Data for Better Lives*. World Bank, 2021. URL: <https://wdr2021.worldbank.org/the-report/>.
- [50] Iris Marion Young. *Responsibility for Justice*. Oxford University Press, 2011. DOI: 10.1093/acprof:oso/9780195392388.001.0001.