Evidential Dependence and Institutional Correction

Flyxion
Independent Researcher

November 7, 2026

Abstract

Institutions rely on chains of grounds, in which reports support findings, findings support determinations, determinations support permissions, and permissions support acts. Two questions about such chains are usually treated separately. The first asks how much support a claim actually has when several apparently separate reports may descend from a common source. The second asks which later judgments, permissions, and commitments must be reconsidered when one of the grounds fails. The first reconstructs the ancestry of a claim, and the second follows the consequences of its revision. This essay argues that they are two directions of a single problem, the preservation of dependence through the transformations that institutions routinely perform. On the first question, the essay holds that the evidential contribution of an additional report depends on what it adds beyond its shared dependencies, so that neither independence nor dependence settles the matter. On the second, it holds that the failure of a premise neither falsifies every downstream conclusion nor requires the reversal of every downstream act, and that a useful record of dependence must therefore preserve the role each ground played. The essay distinguishes three operations that a defective finding can call for, namely correcting the evidence, revising a status, and undoing an act, and argues that they are different operations with different conditions. A hypothetical case involving the certification of a structural fastener runs through the argument.

1 Introduction: Two Directions of Dependence

A regulator certifies a structural fastener for use in occupied buildings. The certification rests on three documents. The first is a manufacturer's test report. The second is a trade association summary that describes the product as having passed rigorous testing. The third is a report from a testing body that applied the same protocol to samples supplied by the manufacturer. A fourth document, a field report from a contractor who used the product on several sites without incident, arrives later and is added to the file. The file

now holds four items, and the regulator's decision memorandum describes the certification as supported by multiple independent sources.

The description invites two questions that point in opposite directions along the chain of grounds. The first asks backward: how much support does the claim that the fastener is safe actually have? The trade summary quotes the manufacturer's report. The testing body's result depends on the manufacturer's samples and the manufacturer's protocol. Whether the contractor's experience adds anything depends on how the sites were chosen, how long the product has been in service, and what kind of failure would have been visible. A count of four documents does not answer the question. The second asks forward: suppose that the manufacturer's test is later found to have used a miscalibrated rig. Which determinations, permits, installations, and contracts rested on it, and what should be done about each? A building has been fitted with the fastener under a permit that cited the certification. A supplier has contracted to deliver more. The question is not whether the original claim was wrong, which may or may not be so, but what follows from the fact that one of its grounds has failed.

The two questions are usually asked in different places. The first belongs to the study of evidence and testimony, where the independence of sources is a central concern. The second belongs to the practice of institutional correction, and to the older study of how a system of beliefs is to be revised when a premise is withdrawn. The thesis of this essay is that they are the two directions of one problem. In each, the object is the structure of dependence among grounds, and the failure in each case is the loss of that structure through the ordinary operations of summarizing, certifying, citing, delegating, and archiving. An institution that has kept the structure can say both what an apparent confirmation adds and what a later defeat withdraws. An institution that has not can say neither, and will tend to overcount its support while it is stable and to overreact or underreact when it is disturbed.

2 Ancestry: What an Additional Report Adds

The first problem is to say how much a further report contributes to the support of a claim. It is tempting to answer by counting, on the principle that more reports are better, and the principle is sound when the reports are independent in the right way. The long-standing observation that the reliability of a collective judgment grows with the number of judges holds only when their errors are not correlated, and the result fails, sometimes dramatically, when judges share sources of error [8]. The same point has been made about testimony: a listener who hears the same claim from several speakers has more reason to believe it only if the speakers have not all drawn it from one place, and the study of expertise has noted

how readily non-independent experts are counted as independent [5].

The correction, however, is easily overstated. A natural summary is that corroboration counts only where sources are independent, and the summary is wrong in both directions. Dependent sources can provide additional support when they contribute distinct observations, checks, or analyses, even though they share some premises or infrastructure. Two testing bodies using the same instrument design are neither wholly independent nor necessarily redundant: they share the risk that the design is flawed, and they do not share the risks of calibration, sample selection, operator error, or local conditions. Conversely, two organizations that publish independently written reports may contribute no additional observational evidence if both repeat the same measurement. The independence of authorship is not the independence of observation. What matters is what each report adds beyond what it shares with the others.

The point can be stated with some precision. For each report R, let $\mathrm{Anc}(R)$ be its ancestry, meaning the observations, instruments, premises, intermediaries, and infrastructure on which it depends, and let $\mathcal{F}(R)$ be the set of ways in which R could be wrong. Ancestry and failure modes are related but not identical: a shared instrument design contributes a failure mode common to every report that uses it, while a shared funding source may contribute none. Consider reports $R_1, \ldots, R_k$ bearing on a claim. The ways in which all of them could be wrong together form the set

$$\mathrm{CM}(R_1, \ldots, R_k) \;=\; \bigcap_{j=1}^{k} \mathcal{F}(R_j),$$

which can be called the set of common-mode failures. Adding a report cannot enlarge this set, since the new report's failure modes are intersected with the existing ones. The contribution of the new report lies in how far it shrinks the set. A report that merely repeats an earlier measurement has the same failure modes as that measurement, and leaves the set unchanged. A report that rests on a separate observation by a different method has failure modes that overlap little with the others, and shrinks the set considerably. A report that uses a shared design but different samples, operators, and conditions falls between the two, removing some common-mode failures and leaving the design-level ones in place.

Definition 1. The *marginal evidential contribution* of a report R_k, relative to reports $R_1, \ldots, R_{k-1}$ already considered, is the reduction it effects in the set of common-mode failures that bear on the claim.

The definition has no numerical content, and none is intended. The sets are not given by nature. They are the product of an assessment, and the assessment can be mistaken. The definition serves to state what the assessment is about, and to show why neither a

count of reports nor a bare verdict of independence answers the question. Institutional multiplication becomes misleading when the number of reports is allowed to substitute for an assessment of what each adds beyond its shared dependencies. The regulator who writes "multiple independent sources" has made a claim about common-mode failures without examining any.

Provenance is the information needed to make the assessment. A body of work on data provenance distinguishes where a piece of data came from from why it appears in a result [2], and both are needed here. Where the trade summary's description of the testing came from is a question about ancestry. Whether the testing body's result has any failure mode that the manufacturer's result lacks is a question about independence of failure. A procedure that counts corroboration without access to either will be unable to tell repeated transmission from separate observation, a deficiency that intelligence analysts have long described under the heading of circular reporting [7].

3 Consequences: What a Defeat Withdraws

The second problem arises when a ground fails. Suppose that the manufacturer's test is found to have used a miscalibrated rig, so that its results cannot be relied on. The tempting inference is that everything built on the test is now unsound, and the tempting remedy is to retract everything downstream. Both are mistaken, for reasons that are best seen by looking at what could depend on a finding.

The failure of a premise does not establish that every downstream conclusion is false. A conclusion may have another sufficient justification. If the testing body's independent result supports the safety of the fastener by itself, then the certification may stand after the manufacturer's report is withdrawn, perhaps with a reduced margin. A conclusion may also hold for reasons unrelated to the defect, since a defective ground can support a true claim. It does not follow either that every resulting act should be reversed. An authorization may have expired in the meantime, so that its later invalidity adds nothing to the case for reconsidering it. An executed act may have created obligations that survive the failure of its original rationale, as when a contract was lawfully made on the basis of a certification later found unsound. And a finding may have been relevant to a decision only as background, so that the decision is unaffected by its failure.

A dependency record that is useful for correction must therefore preserve the role of each dependency, and four roles seem the minimum. A finding may have supplied *necessary* evidence, so that the conclusion cannot be maintained without it unless another route is found. It may have supplied *one sufficient justification among several*, so that the conclusion survives if another remains. It may have supplied a *condition of authorization*, in which case

the warrant for the act lapses when the condition fails, although acts already performed under the warrant need separate assessment. Or it may have supplied *contextual information* that informed the decision without bearing on its validity, in which case the failure calls for review and not for revision. The failure of a ground has a different consequence in each case, and a record that does not distinguish them cannot say which consequence applies.

Two other familiar forms of record are insufficient. A bibliography lists what was consulted, and says nothing about what the conclusion depended on, since a source may have been consulted and set aside. A chronological audit log records what happened and when, which is valuable for many purposes and does not distinguish between a step that bore weight and one that did not. A record of dependence has to be of a different kind, because correction needs to recover the relations among the items and not their order.

This need has an old and instructive counterpart in the study of artificial reasoning. Systems built to maintain a set of beliefs record, for each belief, the grounds on which it is held, so that when a ground is withdrawn the beliefs that depended on it can be reassessed [4, 3]. This approach differs from the abstract theory of belief revision, which specifies what a rational agent should retain when forced to give something up and does not record why anything is believed [1], and the distinction between the two approaches has been debated as one between keeping track of the foundations of belief and attending only to its coherence [6]. The institutional case is on the side of foundations: correction requires knowing what rested on what. It also shows the limits of the belief-maintenance model, which is taken up below.

4 Three Operations

A defective finding can call for three different operations, and treating them as one is a source of confusion in institutional correction.

The first is the correction of evidence. The manufacturer's test report is to be withdrawn or annotated, and the claim that the fastener is safe is to be reassessed on the grounds that remain. This operation concerns the epistemic standing of the claim, and its outcome may be that the claim is still well supported, supported less strongly, or unsupported.

The second is the revision of a status. The regulator's certification was an institutional determination, and its status is a matter of the institution's rules. A certification can be suspended pending review, narrowed to certain uses, left in place with a condition, or withdrawn. The decision among these is not settled by the epistemic result. A claim may be supported only weakly and the certification still retained under a rule that tolerates the weakness, and a claim may be well supported and the certification still suspended because the procedure under which it was issued was defective. The relation between the two is

the same as that examined in other contexts between the truth of a determination and its recognized institutional status.

The third is the undoing of an act. Installations were performed, and some were performed under permits that cited the certification. The question is whether they should be removed, inspected, or left in place. That question depends on the evidential position, on the status of the certification, and on matters that neither settles, such as the reversibility of the installation, the risk posed by leaving it, the reliance of occupants, and the obligations that the parties have assumed. Two installations may be physically identical, one performed under a valid permit and one under a permit that has since lapsed, and still stand differently in respect of what must be done, because they are different acts for the institution whose rules govern them.

The distinctions have a practical consequence. A procedure for correction that proceeds directly from the discovery of a defect to the reversal of acts skips two operations, and will tend to produce either excessive reversal or neglect of the survivors of the defect. A procedure that stops at the correction of evidence will leave statuses and acts as they were. Each operation requires its own warrant, and each has to be assessed on its own conditions.

5 Procedural Status and Truth

Two further distinctions mark the limits of what a dependency record can deliver, and they explain why the three operations remain separate even when they begin from the same defect.

Institutions often adopt rules under which a procedural status arises from the passage of time or from silence. A certification is deemed granted if no objection is lodged within a period. A claim is deemed accepted if no party contests it by a deadline. A rule of this kind is constitutive in the sense that it makes something count as having a status within the institution [9]. It does not establish that the proposition accepted is true. When a defect in a ground is found, a status that arose by default has no ground to withdraw in the ordinary sense, because it was never based on the finding. Its revision depends on the rule that produced it and on the conditions the rule specified. A dependency record that treats such a status as resting on evidence would misdescribe it, and the correction that followed would address the wrong thing.

The converse case is the physical transition that becomes an institutional act through the rules, the actor's standing, and the circumstances of its performance. An installation counts, for the institution, as a compliant installation, a non-compliant installation, or an installation of uncertain standing, and the classification depends on matters other than the physical result. Institutional recognition and validity may also diverge: an act may

be recognized, entered in the record and relied upon, while being invalid for want of a condition, and the opposite may occur. Correction has to deal with both the recognized status and the underlying validity, and the dependency record has to preserve enough about the circumstances of performance for both to be assessed.

These qualifications explain why correcting evidence, revising a status, and undoing an act are different operations even when they originate in the same defective finding. A single failure can propagate along several kinds of dependence at once, epistemic, normative, and practical, and the record has to say which kind each link is.

6 Transformations That Lose Dependence

Institutions preserve or lose dependence through the operations they routinely perform, and the losses explain why the two problems are so seldom answered when they arise.

Summarizing compresses a report into a conclusion, and the compression typically removes the conditions under which the conclusion was reached. A trade association's description of testing as rigorous retains the verdict and drops the protocol, the samples, and the limits. Citation points to a source and often fails to distinguish between reliance on the source for its findings, for its method, or for its standing. Certification converts a determination into a reusable token, which is its purpose, and the token does not carry the ancestry of the determination. Delegation passes a task along with a scope of authority, and the grounds on which the delegator acted are usually not passed along. Archiving preserves documents and rarely preserves the relations among them. Translation into another institution's vocabulary maps terms across contexts, and may map a result obtained under one standard into one that the second institution takes to be obtained under its own.

Each of these operations is useful, and none needs to be eliminated. The difficulty is that they routinely discard exactly the information that the two problems require. A summary of a report may be accurate and still not show whether the report has failure modes in common with another. A certificate may be valid and still not show what it would take to review the determination it embodies. Systems built on automated summarization and synthesis, which produce fluent conclusions from many documents, extend the problem, since they can present an apparent convergence of sources without recording that the sources share an origin.

7 What a Record of Dependence Must Preserve

A record adequate to both problems has to preserve, for each claim, determination, authorization, or commitment that it covers, enough to answer a small set of questions. What

are the grounds, and from what ancestry do they come, so that shared origins can be identified? What role did each ground play, in the sense of the four roles described above? Under what conditions and for what period was each ground taken to hold? What warrant carried the step from ground to conclusion, and who held the standing to apply it? These requirements describe a structure, and they can be met in many ways. A dependency graph whose nodes are observations, claims, determinations, authorizations, and commitments, and whose edges are labelled with the roles of the dependencies, is one such way. Institutions have other ways, from careful narrative memoranda to structured registers, and none is privileged.

The record does not need to be exhaustive, and a principle of proportionality governs its depth. A routine and reversible decision that a few people rely upon can carry a compact standing record, and a consequential and hard-to-reverse undertaking, upon which many have relied, calls for a fuller one. The cost of the record is real, and the benefit appears only when something goes wrong, so the temptation to economize is constant. The measure of whether the economy is justified is the cost of the correction that its absence will make impossible.

It should be added that the record is not a validation. A well-kept record of dependence on an unsound ground shows plainly that the conclusion rests on an unsound ground, which is its purpose, and it does not make the ground sound. The distinction matters because records are sometimes treated as though their existence conferred legitimacy on the determinations they describe.

8 Objections and Limits

It may be objected that full dependency tracking is impracticable in large institutions. The objection is correct for the exhaustive case, and the proportionality principle answers it. Most claims need not be tracked in detail, and many that are tracked need only a rough record. The argument is that the structure of dependence cannot be reconstructed after the fact when it was never kept, and that the cases in which reconstruction is most needed, those involving consequential and widely relied-upon determinations, are the cases in which a record is cheapest relative to its value.

A second objection is that roles and independence are matters of judgment, and that records of them will reflect the judgments of those who keep them. This is true, and it limits what a record can deliver. A record of dependence shows what the institution took the dependencies to be, and it can be wrong about them, as the original assessment of independence can be wrong. The record makes the judgment inspectable and thereby correctable. It does not remove the need for judgment about whether the roles were correctly

assigned, and a later reviewer may revise them.

A third objection concerns the belief-maintenance model that the essay has drawn upon. In systems of that kind, the withdrawal of a ground leads to the retraction of every conclusion that depended on it and had no other support, which is the correct behaviour for a set of beliefs. It is incorrect for an institution, where an executed act has produced obligations and reliance that survive the failure of its ground, and where the right response to a defeat is often to preserve what remains valid. The model is thus useful for the structure of the record and misleading for the policy of correction. The institutional case needs the foundations of belief together with a normative account of what should be done about the consequences of its withdrawal, and the two must not be confused.

A fourth objection is that attention to shared dependencies may be used to dismiss corroboration unfairly, since almost any two reports share something. The framework does not support that use. Shared dependencies are often benign and often useful: shared standards make comparison possible, and shared instruments make results commensurable. The question in every case is whether the shared element is a source of common-mode failure for the claim at issue, and that is an empirical and contextual question that cannot be answered by pointing to the existence of a common element.

9 Conclusion

The regulator's file held four documents, and the memorandum called them multiple independent sources. After the argument, the description invites a more specific set of questions. Which of the four documents descend from one observation, and which bring failure modes that the others lack? If the manufacturer's test fails, which of the later determinations leaned on it as necessary, which as one among several sufficient grounds, which as a condition of the permit, and which only as background? The certification must then be assessed as a status, the evidence as a claim, and the installations as acts, each on its own conditions.

The two problems with which the essay began are the backward and the forward direction of a single structure. An institution that preserves the structure can say what an apparent confirmation adds, and can say what a later defeat withdraws. An institution that loses it in the course of its ordinary operations will count its support by the number of documents, and will respond to a defect by reversing too much or too little. The remedy lies in keeping enough of the structure of justification to answer both questions when they arise, in proportion to what depends on the answer.

References

[1] Alchourrón, Carlos E., Peter Gärdenfors, and David Makinson. "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions." *The Journal of Symbolic Logic* 50, no. 2 (1985): 510–530.

[2] Buneman, Peter, Sanjeev Khanna, and Wang-Chiew Tan. "Why and Where: A Characterization of Data Provenance." In *Database Theory, ICDT 2001*, Lecture Notes in Computer Science 1973, 316–330. Berlin: Springer, 2001.

[3] de Kleer, Johan. "An Assumption-Based TMS." *Artificial Intelligence* 28, no. 2 (1986): 127–162.

[4] Doyle, Jon. "A Truth Maintenance System." *Artificial Intelligence* 12, no. 3 (1979): 231–272.

[5] Goldman, Alvin I. "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63, no. 1 (2001): 85–110.

[6] Harman, Gilbert. *Change in View: Principles of Reasoning.* Cambridge, MA: MIT Press, 1986.

[7] Heuer, Richards J., Jr. *Psychology of Intelligence Analysis.* Washington, DC: Center for the Study of Intelligence, Central Intelligence Agency, 1999.

[8] Ladha, Krishna K. "The Condorcet Jury Theorem, Free Speech, and Correlated Votes." *American Journal of Political Science* 36, no. 3 (1992): 617–634.

[9] Searle, John R. *The Construction of Social Reality.* New York: Free Press, 1995.