Warranted Passage: Capability, Authority, Evidence, and Commitment in Institutional Action

Flyxion
Independent Researcher

November 7, 2026

Abstract

This essay develops a theory of institutional failure centred on unwarranted passage between different kinds of judgment. Establishing a fact, possessing a capability, holding authority, and undertaking a commitment are distinct achievements. Each can contribute to the justification of another, but only through an applicable warrant whose conditions have been satisfied. The central concept, relational substitution, names the error of treating a result achieved under one relation as sufficient for a further judgment when the connecting warrant is unavailable, inapplicable, or conditionally unsatisfied. The argument proceeds through a recurring hypothetical case, a reliably identified structural defect in an occupied building, and introduces a small formal vocabulary only after the reader has met the problems it is meant to clarify. Applications to science, security, organizational governance, and artificial intelligence test whether the distinctions survive outside the initial setting. The essay does not claim that one formal scheme settles domain-specific questions. It claims that institutions can contain accurate information, competent participants, and individually valid procedures while producing a decision that their combination does not warrant, and that understanding this possibility requires attention to the passages between achievements.

1 Introduction: The Sentence That Has Already Decided

Consider an ordinary institutional statement: "The evidence establishes that we may proceed." The sentence reads as a report, yet it has already performed a considerable amount of work. Evidence may establish that a building contains a dangerous defect. It may establish that a particular repair would address that defect, that a named person holds relevant authority, or that a standing emergency procedure applies to the situation. These are different conclusions, and establishing one does not automatically establish the others. The sentence compresses several determinations into a single declaration of readiness, and it does so in a

grammatical form that gives a listener no means of recovering which determinations were made, by whom, and under what conditions.

A concrete case will carry the argument. An inspection correctly identifies deterioration in a structural member of an occupied building. The finding is reliable, the danger is material, and a repair method is available. Nevertheless, a series of further questions remains open: whether access to the occupied premises can be had, who holds the authority to approve expenditure, how occupants are to be displaced and for how long, whether a contractor with the required competence is available, and under what conditions work can begin safely. The proposed repair may be entirely appropriate. The justification offered for commencing it may nonetheless be incomplete, and the incompleteness is invisible so long as the strength of the original finding is allowed to stand in for everything that has not yet been settled.

The essay turns on a distinction between the correctness of a conclusion and the adequacy of the passage used to reach it. A decision can happen to be right even though its institutional justification is defective: the repair is begun without the proper approval, the building is saved, and the defect in the passage goes unnoticed because the outcome was good. Conversely, an appropriately authorized and carefully supported decision can produce an adverse outcome, because the world declined to cooperate with a sound procedure. The subject here is the relation that warrants a determination or undertaking, and the argument does not amount to a retrospective classification of every successful outcome as justified and every unsuccessful outcome as mistaken. Outcomes are evidence about passages, and often weak evidence, since a defective passage can succeed and a sound one can fail.

The thesis is that institutions must preserve the difference between what has been established and what remains required for the next undertaking. Their failures can arise from insufficient evidence, from defective rules, or from poor execution, and these familiar sources receive extensive attention elsewhere. A further source arises when a successful result under one relation is substituted for a result that requires another. In such cases no single participant need have erred. The inspector reported accurately, the manager acted in good faith, the contractor was competent, and the combination was still unwarranted. The sections that follow define the relevant relations, introduce a minimal notation for stating the problem precisely, develop the repair case in detail, and then examine how the same structure appears in scientific practice, computer security, organizational governance, and the design of artificial intelligence systems.

2 The Scope of an Achievement

A result has a limited scope. An inspection establishes something about a structure under specified observational conditions: which members were examined, by what method, at what time, with what instruments, and subject to what limitations of access. A credential establishes something about an identity or an entitlement under a verification procedure that checks certain attributes and leaves others unchecked. A delegation establishes something about an actor's authority within a particular jurisdiction, for particular operations, for a particular period. None of these achievements has unlimited consequences, and the habit of speaking of them as though they did is among the principal sources of institutional error.

The institutional importance of a result lies partly in what follows from it and partly in what does not. A finding of structural danger may justify further investigation, satisfy a reporting requirement, or activate an emergency procedure. Which of these consequences actually follows depends on additional rules and circumstances, and the same finding can support different next steps for an engineer, an owner, an occupant, and an emergency official. The engineer may be required to recommend, the owner to notify, the occupant to leave, and the official to inspect further or to order closure. A single finding thus branches into a family of dependent determinations, and what each branch requires beyond the finding itself is a matter of the rules governing that branch.

The truth of a result must also be distinguished from its recognized institutional status. Something can be true without being established in the relevant proceeding: the defect exists whether or not any inspector has found it, and a fact that no procedure has certified cannot support an institutional act that depends on certification. Equally, a procedurally accepted determination can later prove false, since certification procedures depend on inputs that may be inaccurate and on rules that need not preserve truth. Both possibilities have to be preserved. If they are not, the account either makes institutions omniscient by definition, so that whatever the institution has accepted counts as true, or it treats institutional recognition as capable of determining every kind of truth, so that nothing can be true until some body has said so.

The term *scope of consequence* will be used descriptively for the conclusions and actions that a result warrants under specified conditions. The scope is relational and contextual. Evidence sufficient to open an inquiry may be insufficient to impose a sanction, and a test result adequate to justify further testing may be inadequate to justify deployment, even though the evidence itself remains unchanged between the two uses. What varies is the relation between the evidence and the purpose, together with the conditions under which that relation holds. A record of the scope of consequence, where one is kept, therefore has to state more than the content of the finding. It has to state the purposes for which the

finding was adequate, and by implication the purposes for which no judgment was made.

3 Four Relations Within One World

The title of this section might suggest four separate realities, one for what can be done, one for what may be done, one for what is known, and one for what has been promised. The essay posits nothing of the kind. Capability, authority, evidence, and commitment are relations operating within a shared historical world, and the figure of four worlds names only the difficulty that institutions encounter when they treat these relations as interchangeable, in the way that a traveller who treats a map of roads, a map of property boundaries, a map of terrain, and a map of intended routes as the same map will arrive at the wrong place through no fault of any single map.

Capability concerns what an actor can execute under prevailing physical, technical, organizational, and resource conditions. It is indexed to circumstances rather than understood as a permanent possession. A contractor may possess the relevant skill while lacking access to the premises, the necessary equipment, sufficient time, or a safe working environment, and in that case the contractor is skilled but not presently capable of the repair in the sense that matters for the decision. Capability so understood is a relation between an actor, an operation, and a set of conditions, and it can lapse when the conditions change without any change in the actor.

Authority concerns the normative standing under which an actor may issue, approve, delegate, or perform an operation. Authority should be distinguished from permission. Permission concerns whether an action is allowed; authority concerns an actor's standing to make particular determinations or to alter normative relations. The distinction follows a familiar line in analytical jurisprudence, where a privilege to act and a power to change legal relations are treated as different kinds of legal position [10]. Someone may be permitted to undertake an action without having authority to permit someone else to undertake it, and someone may hold the power to authorize an action that they are not themselves permitted to perform.

Evidence concerns the grounds available for establishing a proposition under an inferential or evaluative procedure. Evidence can concern capability, authority, events, identities, or commitments, and it therefore crosses the other domains without becoming identical to any of them. A record that a delegation was signed is evidence about authority; a record that a contractor completed similar work is evidence about capability; neither is the thing it is evidence of. This crossing is what makes substitution tempting, since the vocabulary of evidence is available everywhere and can be applied to any question without apparent strain.

Commitment concerns an attributable undertaking that enters the history and can change obligations, expectations, reliance, and subsequent options. Several distinct events must be kept apart: considering an action, recommending it, approving it, promising it, and commencing it. Institutions often use similar language for these different events, so that a statement such as "the repair is going ahead" can report a hope, a recommendation, a budgetary approval, a signed contract, or the arrival of a crew, and participants may be uncertain about whether an undertaking has actually occurred and who made it.

These four relations interact extensively. Evidence bears on capability and authority, authority governs which commitments are effective, commitments change what is feasible and what evidence will later be needed, and capability constrains what can responsibly be committed. The purpose of distinguishing them is to make their interactions assessable. A treatment that merged them would be unable to say how one supports another, and a treatment that separated them without connection would be unable to explain why they ever do.

4 A Minimal Formal Vocabulary

A small notation allows the problem to be stated with some precision. The notation is introduced after the problems it addresses have been described, and it claims no more than clarity of statement. Let h be a history, understood as the record of what has occurred up to a point, let i be an actor, and let

$$A_i(h)$$

be the set of candidate actions for that actor at that history. The candidate set contains actions under consideration, including actions that may ultimately prove infeasible or impermissible. Membership expresses relevance to deliberation rather than readiness for execution.

Feasible actions and permitted actions are defined separately:

$$F_i(h) \subseteq A_i(h), \qquad P_i(h) \subseteq A_i(h).$$

Neither set generally contains the other. An actor can be technically capable of a prohibited operation, and an institution can permit an operation that prevailing conditions make impossible. The intersection

$$C_i(h) = F_i(h) \cap P_i(h)$$

contains the actions that satisfy both conditions. Membership in $C_i(h)$ does not establish that an action has been selected, adequately evidenced, or undertaken. Feasibility and permission jointly leave room for prudential judgment, for competing duties, for the allocation

of scarce resources, and for refusal. The set marks where action is not excluded, and says nothing about where it is required.

Available evidence is written $E_i(h)$, and procedural derivability is written

$$E_i(h) \vdash_{\mathcal{R}} \phi.$$

The expression states that procedure $\mathcal{R}$, applied to the evidence available to actor i, supports or derives ϕ. The precise force of the turnstile must be declared for each application. Deductive proof, defeasible institutional acceptance, and statistical support are different relations, and treating them as equivalent would reproduce at the level of notation the very conflation that the essay opposes. Procedural certification is also separate from satisfaction in the relevant world or model, written

$$E_i(h) \vdash_{\mathcal{R}} \phi \quad \text{and} \quad M_h \models \phi.$$

A soundness condition would connect the two within a specified domain, but the existence of a certificate does not itself establish that the soundness condition holds. Procedures may depend on inaccurate inputs, on restricted assumptions, or on rules that do not preserve truth.

Authority judgments take the form

$$\mathrm{Auth}(i, \alpha, h),$$

where α specifies the operation over which authority is claimed. The second argument prevents "authorized" from functioning as an unrestricted property of a person. Finally, a commitment event is represented as a transition between histories:

$$h \xrightarrow{\mathrm{commit}_i(a,\Gamma)} h',$$

where Γ specifies the content and conditions of the undertaking. The occurrence of the event is kept separate from its validity or effectiveness. This separation follows the older observation that a purported act of undertaking can fail to achieve its effect while still having been performed [1]: an attempted or unauthorized undertaking can enter the record and produce consequences, including reliance by others, without acquiring legitimate force. A framework that identified occurrence with validity would be unable to describe the most troublesome cases, those in which something has happened that should not have counted.

5 Relational Substitution

The central concept can now be defined.

Definition 1. *Relational substitution* occurs when a judgment established under one relation is treated as sufficient for a further judgment, although the required connecting warrant is unavailable, inapplicable, or has unsatisfied conditions.

A warranted passage is represented schematically as

$$\frac{J_1 \qquad W_{1\to 2} \qquad \text{Conditions}(W_{1\to 2}, i, h, a)}{J_2},$$

where J_1 is the source judgment, J_2 is the proposed destination judgment, and $W_{1\to 2}$ supplies the connection. The conditions are those under which the warrant applies to actor i and action a at history h. The use of the word warrant is deliberate and follows an established tradition in the study of argument, in which the step from grounds to a conclusion is licensed by a general principle whose applicability is itself a further question [19]. The form is kept general because different passages require different kinds of warrants.

Three kinds deserve separate mention. Inferential warrants license a step in reasoning, as when a logical rule or a validated statistical procedure carries a finding to a consequence. Institutional powers license a change in normative position, as when a delegation confers authority or a valid undertaking alters obligations. Practical procedures license, or more exactly produce, a material effect, as when an execution procedure carries an approved plan into an altered building. Calling all of these bridges must not imply that they operate in the same manner. A logical rule does not create anything; a delegation creates a standing that did not exist; an execution procedure changes the world. A constitutive account of institutional facts, on which certain things count as certain other things within a context [16], helps explain why the second kind cannot be replaced by inference alone: that a signature counts as approval is a fact about the institution, and no amount of evidence about the signature will make it so if the institutional rule is absent.

Absent warrants must be distinguished from unspoken ones. Routine institutional activity depends on standing procedures that participants need not recite on every occasion, and a passage can therefore be justified without being explicitly narrated. A foreman who begins routine maintenance does not cite the delegation that permits it, and the omission is not a defect. Conversely, a written justification can cite an expired delegation or an irrelevant rule and remain invalid despite its apparent completeness. Inspectability and validity are related but distinct properties of a passage. An inspectable passage can be examined, and the examination may show it to be invalid; an uninspected passage may be perfectly sound. The relation between the two matters chiefly when disputes arise, and the later

discussion of records returns to it.

Warrant chains require a declared stopping point. An operational assessment can take a governing rule as given while a constitutional, ethical, or political inquiry questions the legitimacy of that rule. In the vocabulary of legal philosophy, the chain of validity must end somewhere in a rule that is accepted rather than derived [9]. The framework developed here makes that boundary visible rather than claiming to settle every foundational dispute. An analysis that traces a decision to a rule of the institution has not shown the rule to be just, and an analysis that questions the rule has not thereby shown that the decision was invalid within the institution.

6 The Repair Case: One Finding, Several Unfinished Questions

Return to the damaged building. The inspector's report establishes a defect with a stated degree of confidence and identifies the limitations of the inspection: which areas could not be reached, which assumptions about material condition were adopted, which tests were not performed. The report may support a conclusion about risk without determining the full scope or timing of repair. A competent reader can extract from it that the risk is real and serious. The report does not say, and was not designed to say, what is to be done, by whom, or on whose decision.

The legitimate possibilities are many. The finding might trigger notification of the owner and of the relevant public body, a second and more intrusive inspection, temporary shoring, restricted access to part of the building, evacuation, or permanent repair. Each proposed action has different feasibility conditions, different authorization requirements, and different consequences for occupants. Temporary shoring may be feasible within hours and may require only the owner's consent; evacuation may require an order from a public official; permanent repair may require permits, funding, and an agreed schedule with tenants. The finding supports all of these as candidates, and none of them follows from it alone.

Examine the sentence "We know the repair is necessary, so the contractor can begin." Its first clause may concern structural necessity, which is a matter of evidence. Its conclusion presupposes access to the premises, authority to approve the work, funding, a safe sequence of operations, and an actual undertaking by the contractor. The sentence is not merely confused. It can be analysed exactly: the missing premises are an access warrant, an approval by an actor with standing over this expenditure, a determination that the work can be sequenced safely given continued occupancy, and a commitment event in which a contractor with the relevant competence takes on the work under stated conditions. Naming the missing premises is more useful than labelling the sentence careless, because it identifies

what must now be supplied and who is positioned to supply it.

An emergency variation sharpens the point. A standing emergency procedure may legitimately shorten or alter the ordinary process. If the member is close to failure and the building is occupied, a designated officer may be empowered to order immediate shoring and closure without the usual notice and consent. The applicability of such a procedure depends on relevant conditions, such as the severity and immediacy of the danger and the actor's designated role. Urgency can change the governing warrant without making warrants unnecessary: the officer acts under a different warrant, whose conditions must still be met, and an officer who cites the emergency procedure when its conditions are absent has performed a substitution of the same kind as the contractor who begins without approval.

The opposite failure must be treated with equal seriousness. Officials may demand additional certification after the relevant conditions have already been satisfied, causing harmful delay. A second review is requested because the first was done by someone unfamiliar, a form is required that duplicates one already filed, or approval is withheld pending a consensus that no rule requires. The framework must diagnose unwarranted obstruction as well as unwarranted intervention. A procedure that cannot recognize sufficient grounds fails just as consequentially as one that accepts insufficient grounds, and in the repair case the cost of obstruction is borne by the occupants of a building whose defect is already known. Demanding warrants and demanding unlimited warrants are different postures, and only the first is defended here.

7 The Principal Forms of Substitution

The forms of substitution are best seen through connected examples rather than a catalogue of slogans, and each begins from an everyday inference that is usually harmless and sometimes decisive.

The first form takes capability for permission. An interface exposes an operation, a key opens a door, or a system accepts a command. Successful access establishes practical possibility under those conditions, and it does not alone establish entitlement. A contractor with a key to the building can enter; whether the contractor may enter, and to do what, depends on a separate warrant. The inference from can to may is encouraged wherever systems are designed so that whatever is possible is presumed to be intended, and it is defeated by any case in which the system's boundaries were drawn for convenience rather than to express who is entitled to do what.

The second form takes authority for factual correctness. A designated decision-maker can possess the power to issue an effective determination while being mistaken about the

underlying facts. An official may validly declare a structure safe, and the declaration has its institutional effect, which may include releasing the owner from certain duties and permitting reoccupation, while the structure remains unsafe. The determination's institutional effect must be separated from its epistemic accuracy. A treatment that merged them would be forced to say either that the declaration made the building safe or that the declaration had no effect, and both are false.

The third form takes evidence for authority. Establishing that a course of action would be beneficial or safe does not itself identify who may approve it. A thorough analysis showing that shoring is the best response does not make the analyst the person who may order it. At the same time, evidence about a delegation can help establish authority, since the existence and scope of a delegation is a matter of record, and this is why the missing warrant must be specified rather than assuming an absolute barrier between domains. The evidence that a delegation exists is not itself the delegation, yet an institution can reasonably rely on such evidence when determining who holds authority.

The fourth form takes permission for obligation or commitment. An actor may be allowed to proceed while retaining legitimate reasons to decline, since permission opens an option and leaves the selection open. A permitted option becomes an undertaking only through an applicable decision or commitment process. The substitution is common in organizations where the phrase "we are cleared to proceed" is heard as "we are going to proceed," and where a later decision to wait is experienced as a reversal rather than as the exercise of a choice that the permission left open.

The fifth form takes commitment for retrospective authorization. An organization may treat the resources already spent or the promises already made as evidence that the undertaking must have been legitimate. If funds have been paid and a crew has been scheduled, the inference runs, someone must have approved the work. The inference reverses the usual direction of warrant, and it can make correcting an unauthorized commitment increasingly difficult as others begin to rely upon it. Each additional party who orders materials, adjusts a schedule, or informs tenants adds to the cost of acknowledging that the first step was never authorized, and the institution is tempted to ratify rather than examine. A commitment that has been made without authority is a fact in the history, and the question of what to do about it is a distinct question from whether it was valid when made.

8 Evidence Has a Scope of Use

The phrase "evidence has a jurisdiction" can serve as a metaphor for purpose, relevance, procedural admissibility, and inferential reach, provided it is not pressed. Evidence is not generally an authority-bearing actor, and its jurisdiction is not a literal legal property. The

metaphor marks the fact that a body of evidence can be adequate for one use and inadequate for another without any change in its content.

The same record can adequately establish one proposition while leaving another unsettled. A photograph may establish visible damage without establishing its cause. A successful test may establish performance under tested conditions without establishing reliability across the conditions of deployment, and a claim about a phenomenon must be distinguished from the data that are used to detect it, since data are produced by particular procedures and the phenomena they bear on are stable features of the world that those data support only under specified assumptions [3]. An authenticated message may establish attribution without establishing the truth of its contents: the message came from the stated sender, and the sender may be wrong or lying. In each case the record's adequacy is relative to a proposition, and a report that claims adequacy without naming the proposition has not said what it has established.

The strength of evidence must also be separated from the institutional permission to obtain or use it. Information can be highly informative while its collection or use violates governing restrictions, as when a sample was taken without consent or a record was obtained outside a lawful process. Conversely, lawfully collected information can be irrelevant or misleading. A procedure that conflated informativeness with admissibility would either exclude useful information on formal grounds when no rule required exclusion, or accept improperly obtained information because of its apparent force. Both strength and admissibility are properties of the evidence in relation to a use, and neither entails the other.

Purpose sensitivity completes the picture. The evidential threshold appropriate to opening an investigation need not equal the threshold for imposing a sanction, authorizing an irreversible intervention, or announcing an established finding. The reasoning is familiar from discussions of inductive risk in science, where the amount of evidence required to accept a hypothesis is held to depend on the consequences of error in the context where the acceptance is used [15, 7]. The same body of evidence can warrant acting provisionally and fail to warrant acting finally. The essay does not adopt any particular theory of how thresholds should be set. It holds only that the differences between thresholds are explained by the consequences and governing standards of the particular determination, and that a threshold which ignores them is arbitrary in one direction or the other.

Preserving evidential scope therefore requires recording what a finding establishes, under which assumptions, for which purpose, and with what unresolved limitations. A record that keeps only the conclusion invites its use for purposes the evidence was never assessed against, and a record that keeps the surrounding conditions permits a later reader to ask whether the finding reaches the question now being put.

9 Distributed Evidence and the Problem of Composition

Suppose three actors, A, B, and C, each hold evidence that is insufficient by itself for a particular determination:

$$E_A \nvdash_{\mathcal{R}} \phi, \qquad E_B \nvdash_{\mathcal{R}} \phi, \qquad E_C \nvdash_{\mathcal{R}} \phi.$$

Suppose also that a suitable combination would be sufficient:

$$E_A \cup E_B \cup E_C \vdash_{\mathcal{R}} \phi.$$

The union notation initially describes available informational material. It does not establish that an authorized and reliable aggregation has occurred. In the repair case, one actor holds the inspection report, another holds records of earlier alterations to the structure, and a third holds the tenants' complaints about movement and cracking. Together these may justify an order that none of the three can justify alone.

Distributed availability differs from operational accessibility. Records can exist across an organization while no actor has access to all of them, no procedure permits their joint use, or no recognized role can issue the resulting determination. The conclusion is in some sense implicit in the organization's holdings, and yet no one is placed to draw it. It would be an error to say that the organization knows the building is unsafe merely because its parts collectively hold the relevant records; the more careful statement is that the information is present and the passage that would convert it into a determination has not been made.

Provenance and dependence complicate the aggregation further. Apparently separate reports can derive from one witness, one database, or one generated account, and aggregation can increase the number of records without increasing independent support. Three records that trace to the same original measurement are one measurement counted thrice. Any procedure that counts corroboration must preserve enough provenance to distinguish repeated transmission from separate observation, and the study of data provenance has developed precise vocabularies for distinguishing where a piece of data came from and why it is in a result [5]. Without such preservation the language of convergence conceals dependence.

Incompatibility between evidential contexts is a further difficulty. Records may concern different times, use different definitions, or depend on mutually inconsistent assumptions. A load rating computed under one set of design standards and a measurement of deflection taken under different loading conditions cannot simply be set side by side. Composition requires establishing that their contents can jointly support the proposed conclusion, and this establishment is itself a judgment that needs a warrant and a responsible person.

The institutional consequence is that individually competent participants can produce a

collectively defective determination when the conditions for composing their outputs remain unestablished. The relevant competence includes an authorized procedure for composition. Organizations that attend carefully to the quality of each contribution and neglect the question of who combines them, and under what rule, will tend to place the most consequential judgment in the one place where nobody has been assigned to make it.

10 Certificates, Delegation, and the Limits of Trust

Institutions use certificates because a certificate allows an actor to rely on a bounded determination without repeating the entire inquiry. A purchaser relies on a certificate of conformity, a builder on a stamped drawing, a clerk on a signed delegation. This reliance is essential to coordination, since no actor can verify everything from first principles, but it also makes the scope and conditions of the certificate consequential. A certificate that is silent about its scope invites reliance beyond what the certifier determined.

A certificate may need to bind several things: the proposition certified, the subject or action concerned, the issuing role, the procedure used, the relevant history or time, and any conditions of validity. These are conceptual requirements to be adapted to each application, and they do not constitute a universal demand for exhaustive documentation. Work on machine-checkable certificates in computing offers a clear illustration, since a proof carried with a program is useful precisely because the consumer can check that the proof is about this program and this safety policy and not some other [13]. The design lesson generalizes: a certificate is useful to a relying party in proportion to how precisely it binds what it certifies.

Verification of issuance must be distinguished from verification of content. Establishing that a recognized authority issued a certificate does not independently establish that its underlying finding is correct. It may nevertheless establish that a particular reliance is procedurally permitted, since institutions often allow reliance on certificates from designated sources whatever the content turns out to be. These are two questions, one about the provenance of the certificate and the other about the truth of what it says, and a relying party who answers the first and believes that the second has been answered has performed a substitution.

Delegation chains raise parallel questions. A valid delegation must concern the relevant operation, remain within the delegator's powers, and satisfy applicable conditions on further delegation. The phrase "approved by an authorized person" conceals these questions unless the scope of authority is preserved: authorized for what, by whom, up to what limit, and with what power to pass the authority on. Distributed systems research has built formal accounts of how authority propagates along chains of delegation and of how much a relying

party can conclude from a chain [12], and the formal lesson is the same as the institutional one, that a chain is only as extensive as its narrowest link.

The tension that ends this discussion is that between efficient reliance and recoverable grounds. Institutions need practical summaries, and a regime that required every reliance to be justified by reconstructing the entire chain of grounds would grind to a halt. Yet they also need a way to inspect the basis of reliance when a dispute, an anomaly, or a change of circumstances makes that necessary. The workable course is to rely on summaries in ordinary operation while keeping the grounds recoverable, so that the ability to inspect is available when it is needed and not demanded when it is not.

11 Historical Change and the Expiration of Grounds

Warrants are indexed to histories, and histories change. The temporal problem can be stated as

$$E_i(h) \vdash_{\mathcal{R}} \text{Permitted}(i, a, h) \quad \not\Rightarrow \quad \text{Permitted}(i, a, h').$$

The earlier determination may have been entirely correct. Its insufficiency at the later history h' can arise from changed circumstances, with no error in the original judgment. A permission properly granted in March is not thereby a permission in June, and the passage of time can alter whether it applies without any change in the reasoning that produced it.

Variations on the repair case make the point concrete. Occupancy may change, so that a floor which was empty when shoring was approved now houses a tenant. Temporary supports may deteriorate. Access consent may be withdrawn by an owner or tenant. The proposed method may be revised, so that the work to be done is no longer the work that was approved. A determination attached to one operation cannot automatically authorize a materially altered operation, and the judgment of when an alteration is material is itself a point at which a responsible actor is needed.

Persistence conditions differ across warrants. Some remain applicable until a specified event, such as the completion of the work or the revocation of a delegation. Others expire at a deadline, as with permits that lapse after a fixed period. Others require fresh verification near the time of execution, as with a safety check whose value diminishes with the interval between the check and the action. The appropriate persistence rule depends on the relation and on the consequences of change. A rule that is adequate for a standing delegation is not adequate for a verification whose subject can alter quickly.

A further distinction concerns the logical form of withdrawn authority. It is tempting to represent revocation as the mere negation of authorization at the present history, writing $\neg\,\text{Auth}(i, \alpha, h')$. The representation discards something that matters. A history in which

authority was held and later withdrawn,

$$\mathrm{Auth}(i, \alpha, h_1) \ \wedge \ \mathrm{Revoked}(i, \alpha, h_2), \qquad h_1 \text{ preceding } h_2,$$

differs from a history in which authority was never conferred, although the present value is the same in both. Acts performed between h_1 and h_2 were performed under authority and may remain valid; acts performed after h_2 were not; and a third class, acts begun before the withdrawal and completed after it, requires a rule that the bare negation cannot supply. The conditions of the revocation itself also bear on the answer. A revocation may take effect when it is issued, when it is communicated, or at a stated later time, and an actor who proceeds in ignorance of a revocation that has been issued but not yet communicated stands differently from one who proceeds after notice. Third parties who relied on the apparent authority of the actor, such as a supplier who accepted an order, have positions that depend on the same history. An institution that records only the current value of each authorization cannot answer any of these questions, and for that reason the history of an authorization has to be treated as part of its content and not merely as a log of how the present value came about.

The same structure appears when a permission is checked at one moment and exercised at another, and the gap between check and use is the point where warrants can silently lapse. The problem has been studied in detail in the security literature on race conditions between the checking of a permission and the performance of the operation [2]. What that literature treats as a technical hazard is an instance of the general phenomenon described here, namely that a warrant has a time of validity and an action has a time of execution, and the two must be brought into agreement.

Later evidence can defeat an earlier justification without making the earlier decision unreasonable under the evidence then available. Accounts of defeasible reasoning distinguish between a defeater that gives a reason to believe the opposite and one that merely removes the connection between the grounds and the conclusion [14], and both types occur in institutional life: new findings may show that the structure is sounder than thought, or may show that the earlier test did not measure what it was taken to measure. Historical evaluation must therefore distinguish what was true, what was available, what was established, and what could reasonably have been undertaken at the relevant time. Judging a past decision by what is now known conflates these, and the conflation produces a hindsight that is unjust to those who acted well and unhelpful to those who must learn from what happened.

The conclusion is that sound institutional passage requires both a warrant and a defensible account of its continued applicability. A warrant that was good when issued is not

for that reason good now, and an institution that does not ask the second question will occasionally act on grounds that have quietly expired.

12 Commitment, Reliance, and Responsibility

Commitment returns as an event that can alter the institutional history. An undertaking can coordinate future activity by giving others grounds to plan, to allocate resources, or to refrain from alternative action. A tenant who has been told that the repair begins on a given date arranges to be elsewhere; a supplier orders materials; a lender releases funds. This coordinating function is the reason institutions need commitments at all, and it explains why the unauthorized or defective commitment is so troublesome: it produces reliance in the same way as a valid one, and the reliance is real whether or not the undertaking was.

The content of a commitment must be distinguished from the assumptions used to justify it. An actor may undertake to perform a repair by a date while relying on an estimate about the delivery of materials. If the estimate fails, the question of what remains owed depends on the terms of the undertaking and the governing rules, and not on the fact that the estimate was reasonable. Some undertakings are conditional on their assumptions, and the failure of an assumption discharges or alters them; others are unconditional, and the actor bears the risk of the estimate. Which kind a given undertaking is must be recoverable from its terms, since the parties' reliance will have been shaped by it.

Responsibility has several dimensions. A person may be answerable for making a determination, for approving an undertaking, for executing work, for monitoring changed conditions, or for correcting an error. These responsibilities can be distributed, and distribution is often sensible, but it must not make every consequential passage attributable to nobody. The problem of many hands in public administration, in which responsibility for a collective outcome is diffused across participants until no one is answerable [18], is the form this takes in institutions, and its general analysis distinguishes several ways that answerability can be structured [4]. The framework developed here suggests a design requirement: for each passage between relations there should be an identifiable role responsible for checking that the warrant applies.

Refusal, deferral, withdrawal, and amendment are possible institutional acts with their own conditions. A framework centred on commitment must accommodate a justified decision to wait and a justified revision of an earlier undertaking, because otherwise it would treat every decision to defer as a failure. An official who declines to order evacuation because the evidence does not yet meet the relevant standard has acted, and has acted within the space that permission leaves open. An official who withdraws an approval on discovering that the funding source has been withdrawn has acted also. These are acts with their own

warrants, and they should be assessed by the same standards as acts of undertaking.

Revocation does not erase history. A commitment may be withdrawn or discharged, yet others may already have relied on it, and the consequences of that reliance persist. Correction therefore concerns both the future status of the undertaking and the consequences already produced. An institution that cancels a repair order has changed what will happen, and it has not changed what has already happened: materials may have been purchased, tenants relocated, a schedule displaced. The responsibility to address those consequences does not disappear with the revocation and may arise from it.

13 Applications: Science, Security, Governance, and AI

The applications test the explanatory reach of the framework while preserving the differences between domains. Each identifies a particular passage, the warrant required for it, and a recognizable way that the warrant can fail.

In science, a phenomenon, a measurement, an inference, an accepted finding, and an institutional decision made using that finding are different items. The passage from measurement to accepted finding requires methods, replication, and the judgment of a community. The passage from accepted finding to decision requires a further warrant that depends on purpose. Results established under limited experimental conditions can be extended beyond their warranted scope, as when an effect demonstrated in a narrow population is applied to a broad one, or when a laboratory finding is treated as a field guarantee. Evidential acceptance must also be distinguished from decisions about publication, funding, clinical application, or regulation, each of which introduces further purposes and responsibilities. A finding can be well established and still not warrant a particular policy, and the distinction is obscured whenever the expression "the science says" is used to cover both.

In security, the relevant items are request attribution, credential verification, policy satisfaction, execution capability, and completion. An authenticated request can still be unauthorized, since knowing who is asking does not settle whether they may. An authorized request can fail technically, since permission does not confer capability. The classic illustration is a program that holds legitimate authority for one purpose and is induced by a caller to apply it for another, so that an operation is performed for a party who lacked the standing to request it [8]. Particular attention belongs to the possibility that the relevant permission changes between checking it and performing the operation, as discussed above. The object-capability tradition adds an instructive complication: it deliberately arranges that possession of an unforgeable reference conveys the authority to use it [6, 11], which means that it makes a controlled and intended identification of capability with authority. The identification is safe there because the design guarantees that references cannot be

forged or acquired by unauthorized parties, and the case shows that substitution is a fault only where the supporting guarantee is missing, and that a system may legitimately collapse two relations when it has the means to keep the collapse honest.

In organizational governance, a recommendation can become a decision through circulation rather than through an identifiable act of approval. A memorandum is passed among departments, each of which notes it and adds a comment, and at some point planning, budgeting, or procurement begins. Participants may infer that someone else has authorized the work because the activity of preparing for it has started, and each participant's inference is reasonable given the others' behavior. The framework exposes where expectation has substituted for an attributable undertaking: there is a point at which no one decided, and yet the organization is acting as though a decision had been made. Studies of organizational decision making have long noted that bounded attention and divided roles produce outcomes that no individual chose [17], and some analyses of well-known institutional failures describe how a sequence of locally conforming judgments, each compliant with the procedures then in use, led to an outcome that the procedures were meant to prevent [20]. The present framework offers a way of locating such cases without assigning blame to a single participant: the passage between the recommendation and the undertaking lacked a warrant that anyone was responsible for supplying.

In artificial intelligence, a system can be epistemically correct and institutionally wrong. A correct answer about safety or efficiency does not establish the requesting actor's authority, the continued validity of permission, or the existence of a commitment to execute. If an assistant correctly reports that a repair is structurally necessary, it has established a fact, and it has not thereby determined whether the person addressing it may order the repair, whether an earlier permission still stands, or whether anyone has undertaken to carry the work out. The difficulty is aggravated by the interface. When factual assessment, recommendation, approval, and action are presented through the same conversational channel, in the same tone and in the same grammatical forms, the distinctions that institutions maintain between them can be obscured, and a user may take a fluent recommendation for an approval or an approval for an execution. The framework suggests that the design of such systems should make the passage between these relations explicit: a system that assesses should say that it assesses, one that recommends should say that it recommends, and one that acts should do so only under a warrant that can be named and checked.

The applications end by resisting the claim that one formal scheme settles all domain-specific questions. The framework identifies the questions that must be answered and the passages that require warrants. The relevant institutions still supply the substantive standards: what counts as adequate evidence of structural danger, what delegation is sufficient to approve expenditure, what verification suffices before execution. A general theory of

passage cannot replace the particular knowledge of the practice in which passages occur, and its use is to direct attention to where that knowledge is required.

14 Restraint Is Not a Default

A familiar engineering rule holds that when a system cannot establish that an action is permitted, it should deny the action. In the notation used here the rule reads

$$\neg\,\mathrm{Cert}_i(a) \;\Rightarrow\; \neg\,\mathrm{commit}_i(a),$$

where $\mathrm{Cert}_i(a)$ states that a recognized procedure has certified the action for actor i. The rule is often called fail-closed, and in many settings it is exactly right. A system guarding credentials, or a procedure that releases irreversible and dangerous materials, should not proceed on a presumption in its own favour. The rule is nonetheless a substitution of the kind this essay has been describing, and its limits are instructive.

What the rule treats as one thing is two. That an action has not been certified is an epistemic and procedural fact: no recognized procedure has yet established that it is permitted. That an action is impermissible is a normative fact: something rules out its performance. The rule moves from the first to the second by stipulation, and the stipulation is a warrant that requires conditions. It is appropriate where wrongful permission is much costlier than wrongful denial, where denial is cheap to reverse, and where the uncertainty is likely to be resolved soon, since waiting is then inexpensive. It is inappropriate where denial is itself a consequential act, where its costs fall on people other than the one who applies the rule, or where the uncertainty will not be resolved by waiting.

The design of exits in buildings provides a plain illustration. A door whose electronic lock stays secured when power or verification fails protects property, and a door whose lock releases in the same circumstances protects the people inside. Neither is the correct default in general, and the choice depends on what is at stake, which is why safety codes impose requirements on egress that a purely access-oriented design would not choose for itself. The repair case displays the same structure. An order to vacate withheld until the hazard has been fully certified is fail-closed with respect to the order, and fail-open with respect to the risk borne by the occupants. Every default of denial is a default of exposure to something else, and a rule that names only the first hides the second.

The deeper question is what relation should hold between epistemic uncertainty and action. The framework suggests that no single relation holds across all domains, and that the answer is to be specified per class of action, in terms of the asymmetry of costs between the two kinds of error, the reversibility of each course, the expected duration of the uncertainty, and the identity of those who bear the cost. Denial must itself be an answerable act: an

official who withholds approval has made a determination that can be attributed, reviewed, and corrected, in the same way as one who grants it. A closed default may therefore be adopted as a standing warrant for a particular class of actions, provided that it is justified for that class. It cannot be adopted as a universal principle without reproducing, in the form of reflexive caution, the same neglect of conditions that reflexive confidence displays. This is the principle of proportionality applied to restraint, and it supports the earlier claim that both action and restraint must be answerable to sufficient grounds.

15 Objections and Limits

It may be objected that the essay merely restates familiar distinctions. The objection has force at the level of the individual distinctions. The separation of permission from power, of evidence from authority, and of inference from undertaking is well established in several literatures, and many institutions have carefully institutionalized them. The contribution is to examine their composition, and especially the conditions under which a valid result is transported into a different kind of determination. Familiar distinctions may be well respected one at a time and neglected at the junctions, and the junctions are where the cases discussed above arise.

A second objection holds that requiring warrants creates endless bureaucracy. The framework does not demand a separate document or approval for every step. Standing procedures, competent delegation, and established presumptions can support efficient passage, and a passage that rests on a standing warrant is warranted even when nothing is written down on the occasion. The practical aim is to preserve material conditions and to assign responsibility for checking them. Bureaucratic excess, which the repair case showed in the form of unwarranted obstruction, is itself a failure of the kind the framework diagnoses, because it demands grounds that no rule requires.

A third objection is that all action occurs under incomplete information. The premise is accepted, and the response is to distinguish a warrant from a guarantee. A procedure can justify action under uncertainty when it specifies how uncertainty, consequences, urgency, and revisability bear on the undertaking. A warrant for shoring a beam on incomplete evidence is a warrant if the procedure that licenses it takes the incompleteness into account, and it is not a guarantee that the beam will hold. Uncertainty enters the account as one of the conditions of the warrant and does not undermine the idea of warrant altogether.

A fourth objection is that authority itself can be illegitimate. Validity within an institutional order must be separated from the moral or political legitimacy of that order. The framework can expose what a decision depends upon without independently vindicating every governing rule. A passage may be valid under the rules and the rules may nevertheless

be unjust, and a critic who reaches the rule has gone past what the framework addresses. The declared stopping points of the warrant chains mark where this inquiry begins, and the essay leaves the inquiry to other treatments.

A fifth objection is that determined actors can manufacture warrants. They can: formal completeness alone cannot establish substantive adequacy, and a well-formed record can describe a passage that was never really justified. A record of the passage makes its grounds inspectable, but inspection still requires judgment about evidence, scope, authority, and the people affected. The framework is a means of locating where such judgment is needed. It is not a procedure that removes the need for it, and it should not be presented as one.

16 Toward a Discipline of Warranted Passage

The constructive consequence is a set of design commitments rather than a procedure. Institutional design should preserve distinctions among the proposition established, the procedure used, the authority invoked, the undertaking made, and the conditions under which the result remains applicable. Where these are fused in a single expression, such as "approved" or "cleared" or "confirmed," the institution loses the ability to ask which of them is meant and whether it has been met.

Records have a role that must not be mistaken for validation. A record should make it possible to reconstruct consequential passages and to identify what remains unsettled. An elaborate history of unsupported claims remains unsupported, and a record that is meticulous about form while silent about warrant gives an impression of rigor without its substance. The test of a good record is whether a later reader, facing an anomaly, can tell which result supported which subsequent act and under what conditions.

A principle of proportionality governs how much explicitness is appropriate.

Principle 1 (Proportionality)**.** *The depth of verification and the explicitness of the warrant should reflect the consequences of error, the reversibility of action, the likelihood of changed conditions, and the extent of reliance.*

Routine reversible operations can use compact standing warrants, and the discipline asks nothing more of them. Consequential undertakings, especially those that are hard to reverse or that others will rely on heavily, may require more substantial justification, including explicit statement of the conditions under which the warrant remains applicable.

Correction connects to institutional competence. A capable institution can recognize that a passage has lost its warrant, suspend or revise the affected undertaking, preserve what remains valid, and address consequences already produced. The ability to do so depends on knowing which result supported which subsequent act, since otherwise an institution

that discovers an error in one finding cannot determine what else rests on it. Recorded dependencies are therefore a condition of correction as well as a condition of review, and an organization that cannot trace the consequences of a bad finding will tend either to ignore it or to distrust everything.

The discipline should be distinguished from indiscriminate caution. The objective is to make both action and restraint answerable to sufficient grounds. An official who delays without ground, or who declines to act when the conditions of an applicable warrant are plainly met, departs from the discipline as clearly as one who acts without it. The measure in each case is whether the passage that was made, or declined, can be traced to grounds adequate for the purpose and the consequences involved.

17 Conclusion: What Remains to Be Established

Return to the opening sentence, "The evidence establishes that we may proceed." After the preceding sections the sentence invites questions that it did not previously seem to raise. What has been established, and by which procedure? For which actor, and concerning which operation? Under what authority, and with what undertaking to follow? At what point in the history does the determination hold, and has anything changed since? The sentence once read as a conclusion now reads as the compression of an unfinished argument.

The repair case resolves through a warranted course of action. The structural finding remains indispensable, since without it no repair would be called for. Proceeding also requires the relevant authority, the practical arrangements for access and sequencing, and an attributable undertaking by a competent party. Where emergency conditions apply, the appropriate emergency warrant can support timely intervention, and the conditions of that warrant must be met on their own terms. Where conditions change between approval and execution, the undertaking must be reassessed within the scope of that change. None of this reduces the weight of the finding. It locates the finding as one achievement among several that the next act depends on.

The central contribution can be restated without claiming a universal explanation of failure. Institutions can contain accurate information, competent participants, and individually valid procedures while producing a decision that their combination does not warrant. Many failures have other sources, and the argument leaves them aside. What it adds is attention to the passages between achievements, where a result earned under one relation is silently asked to do the work of another.

The governing question can be left with the reader in its simplest form: what remains to be established before this result warrants the next undertaking? Truth matters, authority matters, capability matters, and commitment matters, but institutional adequacy depends

on preserving their differences while making their connections answerable to examination.

References

[1] Austin, J. L. *How to Do Things with Words*. Oxford: Clarendon Press, 1962.

[2] Bishop, Matt, and Michael Dilger. "Checking for Race Conditions in File Accesses." *Computing Systems* 9, no. 2 (1996): 131–152.

[3] Bogen, James, and James Woodward. "Saving the Phenomena." *The Philosophical Review* 97, no. 3 (1988): 303–352.

[4] Bovens, Mark. *The Quest for Responsibility: Accountability and Citizenship in Complex Organisations*. Cambridge: Cambridge University Press, 1998.

[5] 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.

[6] Dennis, Jack B., and Earl C. Van Horn. "Programming Semantics for Multiprogrammed Computations." *Communications of the ACM* 9, no. 3 (1966): 143–155.

[7] Douglas, Heather. "Inductive Risk and Values in Science." *Philosophy of Science* 67, no. 4 (2000): 559–579.

[8] Hardy, Norm. "The Confused Deputy: (or Why Capabilities Might Have Been Invented)." *ACM SIGOPS Operating Systems Review* 22, no. 4 (1988): 36–38.

[9] Hart, H. L. A. *The Concept of Law*. Oxford: Clarendon Press, 1961.

[10] Hohfeld, Wesley Newcomb. "Some Fundamental Legal Conceptions as Applied in Judicial Reasoning." *Yale Law Journal* 23, no. 1 (1913): 16–59.

[11] Lampson, Butler W. "Protection." In *Proceedings of the Fifth Princeton Conference on Information Sciences and Systems*, 437–443. Princeton, 1971.

[12] Lampson, Butler, Martin Abadi, Michael Burrows, and Edward Wobber. "Authentication in Distributed Systems: Theory and Practice." *ACM Transactions on Computer Systems* 10, no. 4 (1992): 265–310.

[13] Necula, George C. "Proof-Carrying Code." In *Proceedings of the 24th ACM Symposium on Principles of Programming Languages*, 106–119. New York: ACM, 1997.

[14] Pollock, John L. "Defeasible Reasoning." *Cognitive Science* 11, no. 4 (1987): 481–518.

[15] Rudner, Richard. "The Scientist Qua Scientist Makes Value Judgments." *Philosophy of Science* 20, no. 1 (1953): 1–6.

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

[17] Simon, Herbert A. *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization*. New York: Macmillan, 1947.

[18] Thompson, Dennis F. "Moral Responsibility of Public Officials: The Problem of Many Hands." *American Political Science Review* 74, no. 4 (1980): 905–916.

[19] Toulmin, Stephen E. *The Uses of Argument*. Cambridge: Cambridge University Press, 1958.

[20] Vaughan, Diane. *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. Chicago: University of Chicago Press, 1996.