

The Conditions of Standing

Generation, Inscription, Evaluation, Authorization,
and Commitment as Distinct Transitions

Flyxion
Independent Researcher

October 2026

Abstract

Five operations are routinely run together in discussions of what a system may rely on: generating a candidate, inscribing it in a shared record, evaluating it, authorizing its use, and committing to it. This document treats them as distinct transitions between system states and defines a relation of standing that says when a proposed transition is entitled to count for a stated consequence, in a stated scope, on stated evidence, under a stated evaluation procedure, given the commitments already in force at a stated time. Five propositions are proved or exhibited by countermodel: occurrence, licensing and warrant are independent; standing is not in general preserved under higher consequence, wider scope or accumulating commitments; the use of a procedure does not confer standing on it; reversibility is neither recoverability nor correctness; and inscription counts underdetermine warranted content. Five cases from different fields apply one analytical procedure, and four short programs make selected claims independently checkable. The aim is consolidation of an existing body of work, not a general theory, and the limits of the framework are stated in the final section.

Scope and Borrowing

The word standing is borrowed from law for its structure and for nothing else. In legal practice, standing is not a verdict on whether a party is right. It is a determination of whether a claim is entitled to be heard for a given purpose, by a given body, with given consequences, and the doctrine that governs it varies by jurisdiction and has been described in terms of injury, causation and redressability in United States federal courts [1]. None of that doctrine is adopted here, and no position is taken on how any court ought to decide any question. What is retained is the shape of the idea: entitlement to count is always entitlement relative to a forum, a purpose and a consequence, and it is distinct from both truth and procedural acceptance.

The same shape recurs where no court is involved. A conjecture produced by a model, a document filed in a registry, a result accepted by a proof checker, a drug approved for a use, and a release that other software depends on are all candidates that have been moved some distance toward counting, and in each case it is possible to ask how far they have moved and who may answer for the distance. Epistemic, computational and institutional senses of standing are defined in parallel below, with separate worked cases, and an argument is permitted to pass from one sense to another only through an explicitly stated bridge.

Mathematical proof, empirical evidence, institutional authority and operational safety are not reduced to one another, and much of the value of the framework lies in keeping them apart while using one vocabulary.

The document does not announce a new theory and does not derive other frameworks from a single theorem. It fixes a vocabulary, states five structural results with their assumptions, applies them to five cases, and says where they stop. Commitment and reversibility have a fuller treatment elsewhere in the same body of work [5], which is cited where it bears on a proposition and is not restated. Four programs in the **examples** directory reproduce selected claims in executable form, and a separate map relates the remaining essays in the corpus to the results given here.

1 Five Transitions

A candidate is anything that might be taken to count: a statement, a proof attempt, a model, a policy, a patch, a decision. Whether a candidate counts is rarely settled by a single event. It is usually moved through a sequence of operations, and the operations are easy to confuse because in ordinary speech one word often covers several of them. A claim is said to have been made, published, accepted, approved and adopted as if these were synonyms, when each names a different change in the state of some shared arrangement. This section separates five such changes. They are called transitions because each takes a system from one state to another, and because what is wrong with a given case is usually located in a transition that was assumed to have occurred, or assumed to follow from another, and did not.

Two notions must be kept apart throughout. A transition *occurs* when it is among the things a system actually does, and this is a fact about the history of the system. A transition is *licensed* when the rules of the system permit it in the state in which it is attempted, and this is a fact about the rules. Neither is the same as a transition being *warranted*, which is the subject of the next section and concerns whether it ought to count for a stated consequence on the evidence available. A system can do what its rules forbid, its rules can permit what no evidence supports, and its rules can forbid what the evidence would amply support. Much confusion in discussions of automated systems and of institutions alike comes from reading one of these off another.

Generation is the production of a candidate. A language model emitting a proof attempt is a generation, and so is a person drafting a clause or a search procedure proposing a configuration. Generation adds a candidate to someone's workspace and does not, by itself, add anything to the shared record. Its preservation condition is therefore modest: generating must not silently alter the shared record, and the candidate must remain distinguishable from anything the record contains. Its answerability condition is likewise modest and concerns provenance, since whoever or whatever produced the candidate should be identifiable if the candidate later matters.

Inscription is the entry of a candidate into a shared record. A document filed in a land registry, a commit pushed to a shared repository, and an entry added to a ledger are inscriptions. What inscription changes is who can see and cite the candidate, and it changes nothing about whether the candidate is correct. Its preservation condition is that the record keep the inscribed item distinguishable from items that have undergone any

further transition, so that the presence of an item in the record is not read as evidence that it has been checked. Its answerability condition falls on whoever inscribes, who must be able to answer for the accuracy of what was entered and the provenance attached to it, though not for its truth.

Evaluation is the application of an evaluation procedure to a candidate on the basis of evidence, producing an assessed outcome that is itself recorded. A proof checker run on a derivation, a referee report, a test suite executed against a patch and an audit of a ledger entry are evaluations. The preservation condition is the most demanding of the five, because an outcome without a record of what was evaluated, against what evidence and by what procedure cannot later be reassessed, and an assessment that cannot be reassessed has been turned into an assertion. The answerability condition falls on whoever applies and maintains the procedure, who must be able to say what the procedure does and does not establish.

Authorization is the permission, by a body with jurisdiction, to rely on a candidate for a stated use. A regulator approving a drug for a specified indication, a court admitting evidence for a specified purpose, and a committee approving a design for production are authorizations. Authorization differs from evaluation in that it is an act of an authority and not an assessment of evidence, and it can be wrong while the evaluation behind it was sound, or right while the evaluation behind it was poor. Its preservation condition is that the use and scope for which permission was given remain attached to the permission, since an authorization stripped of its scope reads as a general warrant. Its answerability condition falls on the authorizing body.

Commitment is a change in a system such that reversing it would require further transitions and others have begun to depend on it. A release tag that downstream projects build against, a ratified statute, a signed contract, and a published result on which later work proceeds are commitments. Commitment changes the state in which later transitions are assessed, which is why it receives a separate treatment of its own [5]. Its preservation condition is that the structure of dependence be recoverable, meaning that it remain possible to say who relies on what. Its answerability condition falls on those who made the commitment and on those who have relied on it in ways that could not have been withdrawn at the time.

Commitments can also arise without anyone deciding. When a package, a record or an interface is relied upon by enough parties, withdrawing it has become costly for reasons that none of them chose individually, and a de facto commitment exists in a state where no commitment transition was ever licensed. This is one reason the five transitions cannot be treated as stages of a single pipeline, even though in well-run systems they often occur in the order listed.

Each of the five can occur without the next, and the next can occur without the previous. A candidate can be generated and never inscribed, inscribed and never evaluated, evaluated and never authorized, authorized and never committed to, and committed to without having been adequately evaluated or ever properly authorized. A system that treats any of these arrows as automatic has, in effect, delegated a decision to the order in which operations happen to run. The distinctions are not offered as a taxonomy for its own sake. They identify the places at which something must be added, whether evidence, a procedure, an authority or a dependency structure, for a candidate to count, and they

make it possible to ask in each case which addition is missing.

It is also possible to say why transitions, and not claims, are the primitive objects of this account. A claim does not gain standing by being a claim. It gains standing when a system is permitted to move it from held to relied upon. A proof gains standing when it moves from candidate to lemma available for further use, a policy when it moves from proposal to binding, and a commit when it moves from branch to shared history. In each case the thing that changes is the permitted use of the item by others, and permitted use is a property of a move made in a particular arrangement, not of the item considered in isolation. Treating claims as the primitive hides this by suggesting that standing could be read off the content of the claim, and the third section shows by countermodel that it cannot be.

Two clarifications limit what is meant. First, private belief and unrecorded thought are not treated as transitions. A person may believe a candidate for any reason or none, and the account is silent on whether they should. The subject is what a shared record or a shared arrangement may treat as established, and for that reason the account concerns shared records and consequences and does not concern minds. Second, the five types are not claimed to be exhaustive. Repair, revocation and delegation are plausible further types, and the final section says why they are not treated here.

2 The Standing Relation

Standing is relative. It is never a free-standing property of a candidate and never a universal permission granted by an authoritative evaluator. A statement can be entitled to count as a hypothesis worth testing without being entitled to count as evidence for anything, a model can be entitled to be used for exploration without being entitled to be deployed, a decision can be procedurally authorized while its consequences remain indefensible, and a result can count within the scope in which it was established and cease to count when it is carried beyond that scope. Each of these observations is familiar, and each is usually stated as a qualification added to a verdict. The proposal here is to make the relativity part of the definition, so that the arguments on which a verdict depends are written down wherever the verdict is.

The definitions require a minimal notion of state. A *state* s of a system comprises a shared record of inscriptions with their provenance, a designation of which items are relied upon and for what, and the set $K(s)$ of commitments in force. An *operation* o is a proposed change. A *transition* is a triple $\tau = (s, o, s')$ in which s' is the state that applying o in s would produce. Transitions are state-indexed on purpose: the same operation applied in two different states yields two different transitions, and the second proposition of the next section shows that these can differ in standing even when nothing about the operation or the evidence differs. A *history* is a sequence of transitions in which each begins in the state where the previous one ended, and a transition *occurs* when it is an element of the history the system actually follows. A *rule system* assigns to each state the set of operations it permits, and a transition is *licensed* when its operation is permitted in its initial state.

Definition 1 (Standing). *Let $\tau = (s, o, s')$ be a transition. Let C be a consequence class, the use or effect the transition would license. Let S be a scope, the set of contexts,*

domains or populations in which the result would be relied upon. Let V be the evidence available, understood as a body of observations, derivations, records and testimony. Let P be an evaluation procedure, a rule that takes a transition together with C , S , V and the commitments $K = K(s)$ and returns acceptance or rejection. Let t be the time of assessment, which fixes the version of V and of P that is available. The transition has standing when P accepts it, which is written

$$\text{Std}(\tau \mid C, S, V, P, K, t).$$

The relation is a predicate evaluated relative to a context. It is not a numerical measure, since nothing in the definition assigns degrees, and it is not required to be computable, since P may be a committee, a court or a practice that no algorithm captures. The absence of numbers is deliberate: a number attached to standing would suggest a common scale on which a proof, a clinical trial and a court ruling could be ranked, and the account is built to deny that such a scale exists.

Evidence V and procedure P are separate arguments from the start, and this has a direct consequence. A transition can fail to have standing because the evidence is inadequate for the consequence, because the procedure is defective or ill-suited to the consequence, because it conflicts with a commitment in K , or because of any combination of these. A procedure applied to excellent evidence can still produce a verdict that does not deserve to count, and excellent procedure applied to poor evidence cannot repair the poverty of the evidence. Keeping the arguments separate makes it possible to say which has failed, and the two kinds of failure call for different remedies, since better evidence can be gathered while a defective procedure must itself be assessed, which is the subject of the third proposition.

The remaining arguments do their own work. The consequence class C matters because the same candidate may deserve to count for a use with modest effects and not for one with severe or irreversible effects. The scope S matters because a result established in some contexts may not extend to others. The commitments K matter because they change what a transition would newly make irreversible, which is why standing attaches to a transition in a particular state and not to a candidate in the abstract. The time t matters because evidence accumulates, procedures are revised, and what had standing at one time may not at another, without anyone having been wrong at the earlier time.

The relation separates three things that are usually run together. *Occurrence* is what the system did. *Licensing* is what its rules permit. *Standing* is whether, given V and P and the rest, the transition ought to count for the stated purpose. Licensing is a feature of a rule system and standing is a feature of a warrant relation, and in a well-designed system the rules license exactly the transitions that would have standing. No system is well designed in this sense for all transitions in all states, and the gaps run in both directions. Rules can license what has no standing, and can fail to license what would have it. The first proposition below makes this precise, and the later sections return to the gap repeatedly because it is where most failures are located.

There are three senses of standing, distinguished by what plays the roles of V , P and K . In the *epistemic* sense, the evidence V is a body of observations and arguments, the procedure P is a method of inquiry with norms of inference, replication and error control, and the commitments K are the claims on which a research community has

already begun to build. In the *computational* sense, V is a set of execution traces, proofs, test results or formal derivations, P is a checker, verifier or test harness whose behavior is specified in advance, and K is the state of deployed systems and their dependencies. In the *institutional* sense, V is a submitted record, testimony or filing, P is a decision procedure specified by rules and exercised by a body with jurisdiction, and K is the body of prior rulings, ratified decisions and contracts.

The three senses are defined in parallel so that none is made the default. A proof that passes a formal checker has computational standing for the claim that a specified derivation is valid under the checker’s rules, and it has epistemic standing for the mathematical claim only to the extent that the formalization corresponds to the claim, which is a separate question. A regulator that accepts a verified computation as part of a filing treats it as evidence V in an institutional procedure, and the regulator’s authorization does not become a computational verdict by that route. An argument may pass from one sense to another only through a stated *bridge*, which says what in one sense stands in for V , P or K in the other and what is lost in the translation. Statements that move silently across such a boundary, saying that something was verified and therefore authorized, or authorized and therefore true, are the standard form of the error the framework is meant to expose.

Four short contrasts show how much the relativity matters. A claim that a compound reduces a symptom in a pilot study has standing as a hypothesis warranting a larger trial and has not the standing of evidence for prescribing it. A model evaluated on a held-out benchmark has standing for exploratory use in a research setting and not, on that evidence alone, for decisions about individuals. A procurement decision adopted by the proper committee under the proper rules has institutional standing as a decision, and whether its consequences are defensible is a question about a different evidence base and a different procedure. A result established for a laboratory population has standing within the scope of that population and loses it when applied to a population on which the procedure was never run. In each case no one has made a mistake of fact, and the difficulty lies wholly in what the verdict was permitted to carry.

Finally, the relation is something to be assessed and not something that can be looked up. There is no general algorithm that, given a transition and a context, returns whether standing holds, since P may itself be contested and V may be incomplete in ways that cannot be detected from inside the procedure. The relation is useful because it names the arguments that any assessment has to supply, and because it makes visible when an assessment has supplied fewer of them than the verdict requires. The next section states what can be said in general about how the relation behaves.

3 Propositions

The five results below are structural. Each is stated with its assumptions, each negative result is supported by an explicit countermodel, and none is offered as a claim about the world beyond the models in which it is proved. Where a result has been checked by exhaustive search over small systems, the corresponding program in the `examples` directory is named. The strength of the results is meant to come from the explicitness of their assumptions and not from the generality of their conclusions, and a reader who finds

an assumption unreasonable for a given domain is invited to see that as the beginning of an analysis of that domain and not as a defect in the result.

Independence of occurrence, licensing and warrant

Proposition 1. *Let a rule system and a standing relation be given for a system. Then (a) a transition can occur without being licensed; (b) a transition can be licensed and not occur; (c) a transition can be licensed without having standing; and (d) a transition can have standing without being licensed. Moreover, in any rule system that requires evaluation as a prerequisite of commitment, every licensed history in which a commitment occurs contains an evaluation before it.*

Proof. Claims (a) and (b) follow from the fact that occurrence is a property of the history and licensing a property of the rules: a history can contain an operation the rules do not permit in its state, and the rules can permit operations that the history never contains. For (c), take a rule system in which commitment has no prerequisites and a candidate that P rejects on the available evidence; the commitment is licensed and, by the definition of standing, lacks it. For (d), take a rule system that requires two independent approvals before commitment and a candidate that P accepts on V after one approval; the commitment has standing and is not licensed. The final claim is conditional on the rules and follows from the definition of licensing: if evaluation is a prerequisite of commitment, a history that commits before evaluating contains a transition that is not licensed. An exhaustive search over the 1024 rule systems in which each of five operations may require any subset of the earlier operations, and over all 3906 histories of length at most five, confirms each part; it finds that in the rule systems requiring evaluation before commitment no licensed history commits without evaluating, while in the rule systems that do not require it there are 72,012 licensed histories that commit without evaluating (`ex4_countermodels.py`). $\square$

The proposition concerns the independence of occurrence and warrant and does not assert an unrestricted independence between every adjacent pair of transition types. That an unauthorized commitment can occur without prior evaluation does not show that commitment is licensed without evaluation, and the final claim of the proposition is there to block that inference. Where a system's own rules require evaluation before commitment, the only ways a commitment can occur without evaluation are that the rules were broken or that the rules were defective in what they required, and these are different failures with different remedies.

Non-preservation of standing

The second result requires an explicit assumption about how standing is assessed, because without one nothing can be said about how it changes. The assumption is deliberately minimal.

Assumption A. The procedure P is fixed and deterministic. A candidate is a function h from a finite set of contexts to outcomes. The evidence V records outcomes observed at some contexts. A consequence class carries a confirmation requirement c , and higher consequence classes carry higher requirements. P accepts h for (C, S) on V if and only

if h agrees with every observation in V that lies in S and at least c contexts in S are observed.

Proposition 2. *Standing is not in general preserved (a) when the consequence class is raised, (b) when the scope is widened, or (c) when commitments accumulate. Parts (a) and (b) hold under Assumption A, and part (c) holds in a model whose procedure rejects any transition that removes an item on which a commitment depends. Failure of preservation is not the same as monotone loss: widening the scope can also produce standing that was absent.*

Proof. For (a), let there be two contexts, let V record outcome 0 at the first context, let h be constantly 0, let S contain both contexts, and let $c = 1$ for the lower class and $c = 2$ for the higher. Then h agrees with all observations and one context is observed, so h has standing at $c = 1$ and lacks it at $c = 2$. For (b), let V record outcome 0 at two contexts, let h take the value 0 at the first and 1 at the second, let S contain only the first and S' contain both, and let $c = 1$. Then h has standing in S and lacks it in S' , because it disagrees with an observation that entered the scope. The gain is shown by the same evidence with h constantly 0 and $c = 2$: in S only one context is observed and in S' two are, so standing is absent in S and present in S' . For (c), take the state-indexed model of the branching case below, in which the operation is a rewrite of history, the evidence is that the final tree is unchanged, and P rejects any rewrite that removes a published commit. In a state with no published commits the rewrite is accepted, and in the state obtained by publishing one of the commits it is rejected, with C , S , V and P unchanged (`ex2.branching.py`). An exhaustive search over four contexts finds the minimal cases for (a) and (b) and finds that raising consequence never produces standing in this model while widening scope sometimes does (`ex4.countermodels.py`). $\square$

The proposition does not say that increasing consequence, widening scope or accumulating commitments logically forces standing to disappear. It says that none of these changes can be assumed harmless, and that any claim of standing made at one setting must be remade at another. The commitments case is the most instructive of the three, because it shows why standing attaches to a proposed transition in a particular state and not to an isolated candidate. In the branching example the candidate operation is identical in both assessments and so is the evidence that it preserves the tree, and a stateless assessment, which looks only at the operation and the evidence, returns the same verdict in every state and is therefore mistaken in at least one. What differs is the set of commitments, and with it what the transition would newly make irreversible. This is the result that most clearly requires transitions as the primitive.

Use does not confer standing on a procedure

Proposition 3. *Let P be an evaluation procedure and let I be a finite set of instances on which P has been applied. Let J be a criterion that is independent of P , such as ground truth or a later outcome. Then there is a procedure P' whose acceptances on I coincide with those of P and which disagrees with P outside I . Consequently the record of P 's acceptances cannot by itself establish that P rather than P' is the appropriate procedure for any consequence that reaches beyond I .*

Proof. Define P' to equal P on I and to accept every candidate outside I . The acceptance records of P and P' over I are identical, so no feature of that record distinguishes them. If J agrees with P on I and disagrees with P somewhere outside I , then P and P' cannot both agree with J there, and the record cannot say which does. An instance of the same form arises when P accepts everything it is shown and every candidate shown so far happened to be true: the record is then indistinguishable from that of a procedure that accepts exactly the true candidates. $\square$

The proposition is stated at this level and no further. It says that the fact that a procedure is used to evaluate candidates does not establish the procedure's own standing, and that standing for the procedure has to come from evidence about how it performs against something other than its own outputs, or from an explicit convention or stipulation that the system adopts knowingly. The question of whether this leads to an infinite regress of evaluators is separate and is treated as a limitation in the final section. A procedure can be assessed by another procedure, and formal systems can work relative to stipulated rules, and institutions can confer authority on procedures by conventions that remain contestable, without any claim that every evaluation requires an unending chain.

Reversibility, recoverability and correctness

Reversibility is used in two senses that should be separated. A transition (s, o, s') is *reversible in the record* when some later operation takes s' to a state whose record is the same set of inscriptions as that of s . It is *recoverable* when, in addition, everything that depended on the transition during the interval has also been returned to what it was in s : the copies taken, the actions based on the item, and the consequences realized.

Proposition 4. (a) *A transition can be reversible in the record and not recoverable.* (b) *The reversibility of a transition does not establish that it has standing for the consequences it has while in force.* (c) *Reversibility can legitimately lower the evaluation that a procedure requires of a transition, but it does not discharge the requirement for the interval in which others rely on the result.*

Proof. For (a), let a state consist of a record R and a set D of dependents' copies, let publication add an item x to R , let each dependent copy x into D when it is in R , and let retraction remove x from R without touching D . Publication followed by retraction returns R to its earlier value and leaves D changed, so the transition is reversible in the record and not recoverable. For (b), suppose x is false and that P would reject its publication for a consequence class that includes reliance by dependents. Publication is reversible and, during the interval before retraction, dependents rely on it, so the transition lacks standing for that class while the ability to retract remains. For (c), the ability to retract changes what a rejection would cost and what a mistaken acceptance would cost, and a procedure may reasonably require less of a transition that can be undone cheaply. The requirement is lowered and not removed, because the interval of reliance is itself a consequence. $\square$

The distinction connects to the treatment of commitment and the preservation of future options given elsewhere in the corpus [5] and is not developed further here. Ordinary cases show its force. A published claim can be retracted without erasing its influence on work done in the interval. A change to software can be reverted without reversing

every action taken while it was active. A history rewritten in a repository can be restored locally from stored objects while collaborators' copies have diverged, which is the reason the documentation of the standard version control tools advises against rewriting history that others have based work on [4]. In each case the formal state space supports reversal and the world has not been returned to where it was.

Inscription counts underdetermine warranted content

If warranted items are a subset of recorded items, then the number of recorded items can grow without bound while the number of warranted items stays fixed, whenever unwarranted inscriptions are permitted. That observation is immediate and is not the result. The substantive claim is informational.

Proposition 5. *Let a record be a finite set of inscriptions, each an observation or a derived item citing other inscriptions, and let standing be assigned by grounded closure: observations have standing, and a derived item has standing if and only if it cites at least one item and every item it cites has standing. There are records with the same number of inscriptions, the same number of observations and the same number of citation links whose sets of items with standing differ in size. Hence no function of these counts determines the warranted content of a record.*

Proof. Let the first record contain observations 1 and 2 and derived items 3 citing 1, 4 citing 1 and 2, 5 citing 3, and 6 citing 4 and 5. All six items have standing. Let the second record contain observations 1 and 2 and derived items 3 citing 4, 4 citing 3 and 1, 5 citing 6, and 6 citing 5 and 2. Each record has six inscriptions, two observations and six citation links. In the second record the derived items cite one another in cycles, none of them is grounded, and only items 1 and 2 have standing. An exhaustive search over all 4096 records of four inscriptions finds 25 classes of records agreeing in these three counts, and finds that in 9 of the classes the members differ in how many items have standing (`ex1_ledgers.py`). $\square$

The result identifies what a count-based view of a record cannot see. Two records can look identical to any summary that counts inscriptions, observations and links and differ completely in what they warrant, because what matters is the dependency structure and its grounding, which counts discard. The practical form of the point is that a growing record is not evidence of a growing body of established claims, and that measures of volume, activity or citation are silent about warrant unless they encode the structure that warrant depends on.

A sixth relationship lies outside the scope of this document. When generating candidates becomes cheap and evaluating them does not, the cost of the gap falls on whoever maintains the shared record, and the resulting asymmetry has an economic form that deserves a separate treatment under its own name. That argument presupposes the definitions given here, since it needs standing to be defined before its cost can be discussed, and it is deferred.

4 Cases

The propositions earn their place only if they identify real recurrences in subjects that otherwise differ, and the differences are as informative as the recurrences. Five cases follow, and each is analyzed by the same procedure so that the vocabulary does the same work every time. The procedure identifies the transition at issue, says what plays the roles of C , S , V , P , K and t , names the proposition the case instantiates, and ends with an account of what does not transfer to other domains. The cases are written for this document, and the essays in the corpus that treat the same subjects are cited as sources and are neither excerpted nor summarized, since their terminologies were developed independently and are not assumed to match.

4.1 Self-distillation

The setting is a model trained on outputs that it, or a proxy for it, has generated. The proxy-teacher variant is analyzed in an essay in the corpus [6]. Three steps are easy to run together and are kept apart here: *selection*, in which some of the generated outputs are kept; *training*, in which the model is refit to what was kept; and *external evaluation*, in which the resulting model is measured against something other than its own outputs. Training on unverified outputs does not mean that no evaluation occurs anywhere in the process, and the absence of an explicit verifier does not by itself establish drift, so the case states one failure mode under stated conditions and shows it without implying that it is general.

The transitions at issue are the inscription of generated outputs into a training record and the commitment of the refit model as the source of the next generation’s outputs. The consequence class C is the use of the next generation as a generator or a deployed model. The scope S is the range of items the model is expected to cover. The evidence V is whatever is independent of the model’s own outputs, here a known truth. The procedure P is a verifier that rejects incorrect items, or a metric that scores the result, and the commitments K are the earlier generations already relied upon. The time t is the generation index. The model in `ex3_selfdistillation.py` is a distribution over thirty correct items, with truth following a 1/rank law, and ten wrong items, refit by counts to one hundred samples per generation for forty generations, averaged over twenty seeds.

Under these conditions, refitting on all samples with no fresh data reduces the proportion of correct items that retain nonzero probability from 1.000 to 0.127. Adding a verifier in the loop, which rejects wrong items before refitting, removes the wrong mass entirely and leaves coverage at 0.140, so the verifier establishes that each retained item is correct and does not establish that the retained items are representative. Adding a tenth as many verified fresh samples from the truth restores coverage to 0.553. Selecting samples by the model’s own confidence reduces coverage to 0.040 and incidentally removes the wrong mass, since the wrong items start with low probability. A fourth condition computes the same metrics each generation without altering anything and reproduces the first condition’s trajectory exactly on every seed, which is the sense in which evaluation that only measures is distinguished from evaluation that changes selection or training. The average wrong mass in the unfiltered condition rises from 0.080 to 0.111, but the simulation was not designed to test amplification of error and this figure is not interpreted as evidence of it.

The case instantiates the fifth proposition. Each item in the verified record is individually warranted and the record as a whole is not a faithful account of what it was meant to cover, so the number and correctness of inscriptions underdetermines whether the record can serve its consequence. It also illustrates the difference between a verdict about an item and a verdict about a body of items, which is a difference of scope in the sense of the second proposition.

What does not transfer is considerable. The model is a categorical distribution refit by counts, and nothing about gradient-based training, regularization, architecture or the mixing of real data follows from it. The failure mode holds under the stated conditions and is not claimed for any particular system. Coverage is defined here relative to a known truth, whereas in practice the independent criterion that would make coverage measurable is typically what is unavailable, and that unavailability is itself a limit on external evaluation and not a detail of the toy.

4.2 Corrective Navigation and the Repair Ledger

The setting is a navigable record in which corrections are themselves entries, as in the corrective-navigation specification in the corpus [7]. A repair of an earlier entry is inscribed like any other, and the question is what gives the repair standing over the entry it repairs, since an inscription is not an evaluation and a later entry is not for that reason a better one.

The transitions at issue are the inscription of a correction and the evaluation that would let it supersede its target. The consequence class C is the replacement of the earlier entry, for readers who reach the region of the record through navigation, by the correction. The scope S is the region of the record that the correction claims to repair. The evidence V is whatever shows both that the target was in error and that the correction fixes the error without introducing another. The procedure P is the process by which corrections are accepted, which may be the same as the process for original entries or stricter. The commitments K include everything already relied upon that was built on the target. The time t matters because a correction accepted early may be revised by evidence available later.

Two propositions apply. The third bears on the acceptance of corrections by a repair procedure. A repair procedure that has accepted many corrections has an acceptance record, and by the third proposition that record cannot show that the procedure repairs, since a procedure that accepts every proposed correction produces the same record as one that accepts only correct ones until a bad correction is tested against something independent. The fifth bears on the ledger as a whole. A record to which corrections are continually appended grows in inscriptions, and whether it grows in warranted content depends on whether each correction is grounded in evidence about the item it replaces or only cites a previous entry that cites another, in the cyclic arrangement the fifth proposition exhibits. A record whose repairs refer only to one another has the appearance of continuous repair and the warrant of none.

What does not transfer is the idea that the same correction procedure suits every domain. A correction to a mathematical derivation is checked against rules of inference and can be settled finally. A correction to an empirical claim is a revision of evidence and is provisional. A correction to an institutional record, such as an erratum to a ruling,

is an act of an authority and is valid by virtue of that authority. A repair ledger that records all three in one format has not thereby made them the same kind of thing, and the standing of each is established differently.

4.3 Economic Authorization without Effective Accountability

The setting is a decision about production or allocation that has been authorized by a collective procedure of the kind proposed under the heading of economic democracy, and the source in the corpus is a response to an interview on that subject [8]. The case concerns structure only and takes no position on whether any particular arrangement achieves what it intends. The transition at issue is authorization, and the claim is that the standing of the authorization and the standing of the consequences claimed for the decision are two different things established on different evidence.

The consequence class C is the allocation's effects on those who bear them. The scope S is the population affected, which may be wider than the population that took part in the procedure. The evidence V , for the authorization, is the record of how the decision was made and whether the procedure was followed, and for the consequences, it is information about what the allocation actually does. The procedure P for the authorization is the collective decision rule, and the procedure for the consequences is whatever verification is applied to outcomes. The commitments K are prior decisions, contracts and dependencies that the new decision inherits. The time t matters because outcomes arrive after authorization.

The first proposition applies directly. A decision can be licensed by the rules of the collective procedure and lack standing for the consequences it is said to have, and a decision can be defensible in its consequences while the procedure that produced it was irregular. The case is an instance of the occurrence of a licensed transition not entailing its warrant, and of the licensing not being the same as the standing for any consequence that the procedure does not examine. The second proposition adds that standing for a population that participated need not extend to a population that did not, so a decision that has institutional standing within its scope may lack it when its effects fall beyond. The structure of the claim is that verification of outcomes is a different operation from authorization of decisions, and a system that carries out the first well has not thereby carried out the second.

What does not transfer is the normative content of the procedure. Democratic legitimacy is not merely a warrant relation, and a procedure may have a value for the participants that is independent of how good its outcomes turn out to be. The framework says that procedural authorization does not establish outcome standing, and it does not say that procedural authorization is merely instrumental. It also makes no empirical claim about how often decisions made by collective procedures succeed or fail.

4.4 Reconstructive Memory

The setting is a store in which retrieving an item regenerates it at the moment of use, and the regenerated item is then recorded or relied upon as though it had been stored and fetched. Experimental work on human recollection has long shown that recalled content is rebuilt from fragments and expectations, and can be altered by the circumstances of recall

[2, 3]. A computational analogue is a memory system that holds an index and summaries and generates the text it returns. The case is written last of the five and checked against the definitions of the second section, because it presents the greatest risk of confusing categories.

The transition at issue is an inscription that carries a label, for instance that the item was retrieved, which misdescribes the transition that took place, which was a generation. This is first a failure of provenance. The record states that one kind of transition occurred when another did, and the first proposition applies, since what occurred and what the record says occurred are different matters. Whether the regenerated item has epistemic standing is a separate question, answered by further evidence. The consequence class C is reliance on the item as a basis for action or testimony. The scope S is the situation in which it is used. The evidence V is corroboration from independent records or sources, and the procedure P depends on the sense: in the epistemic sense a comparison with corroboration, in the computational sense a check of the regenerated text against stored source if there is one, and in the institutional sense the rules of a body that admits testimony. The commitments K are the actions already taken on the item, and the time t is the time of recall, which matters because reconstruction can differ across recalls.

Three assessments are kept apart. The mislabel is a provenance failure and is cured by correcting the label, which does not touch the content. The epistemic standing of the content depends on corroboration and may be high or low, and that reconstruction is compatible with accuracy is not in doubt. A computational check that the regenerated text matches a stored source establishes a match and not the truth of the source. An institutional decision to admit the item as testimony establishes admission and not correctness. A failure in one of these is not a failure in another, and presenting all four as one failure, whether of verification or of authorization, loses what the case teaches.

What does not transfer is any estimate of how reliable reconstructed items are. The literature on human recollection is varied, and it is cited only to show that regenerated content can be inscribed as retrieved. Nothing here implies that reconstructed recollections lack standing in general, or that stored items have it in general: in each direction the standing has to be established, not presumed.

4.5 Irreversible Commitment in a Branching System

The setting is distributed version control, in which a history is a graph of commits and branches are named pointers into it. The branching-coordination essay in the corpus treats the setting at length [9], and the case here concerns one operation in it. The transition at issue is a rewrite of history that squashes several commits into one, and the commitment at issue is the publication of some of those commits so that collaborators may base work on them.

The consequence class C is the continued use of the shared history by collaborators. The scope S is the set of clones and branches that contain the affected commits. The evidence V is that the final tree is unchanged by the rewrite, and the procedure P accepts a rewrite when that holds and when no commit that is published is removed. The commitments K are the published commits, and the time t is the time of assessment. The program `ex2_branching.py` evaluates the same rewrite, with the same evidence, in five states that differ only in which commits are published. The rewrite is accepted when nothing

is published or only the base commit is, and is rejected when any removed commit is published. A stateless assessment, which looks at the operation and the tree evidence alone, accepts in all five, and so is mistaken in three. Publishing a commit is a commitment transition, and after it the rewrite that was acceptable is not.

The case instantiates the second and fourth propositions together. By the second, the commitment changes what the rewrite would newly make irreversible, which is why standing attaches to the transition in its state. By the fourth, the rewrite is reversible in the record, since the old objects remain stored and can be restored by their holder, and is not recoverable, since collaborators' copies have diverged and must be repaired by each of them. The standard documentation of the tools advises against rewriting history that others have based work on for exactly this reason [4]. The first section's remark on *de facto* commitment also applies, since a commit that collaborators have begun to rely on is published in the relevant sense whether or not anyone designated it so.

What does not transfer is the clarity of the dependency structure. In version control the dependencies are explicit, because commits are identified by content and clones can be enumerated in principle, so what is published is in principle observable. In most other domains who relies on what is not recorded and has to be reconstructed, and much of the difficulty of commitment lies in that reconstruction. The rule against rewriting published history is also a convention with exceptions, such as private branches and coordinated rewrites that all collaborators have agreed to, and its force is relative to what a given community has committed to and not an absolute property of the operation.

5 Limitations

The framework is a vocabulary and a set of structural relationships, and it is deliberately incomplete. The first limitation is that there is no general algorithm that decides every standing question. The relation depends on an evaluation procedure that may itself be defective, contested or undefined for the transition at hand, and on evidence whose completeness cannot in general be checked from within the procedure that uses it. The relation is a predicate to be assessed, with the arguments that any assessment must supply made explicit, and it is not a function that can be computed once and consulted.

The second limitation concerns what standing cannot remove. It does not eliminate judgment, since choosing a consequence class, a scope or an evaluation procedure is a judgment, and the relation inherits whatever is contestable in those choices. It does not eliminate uncertainty, since a transition can have standing on the available evidence and be mistaken. It does not eliminate institutional conflict, because two bodies with different evidence, procedures and commitments can assign different standing to the same transition without either being mistaken about the facts. A court and a laboratory, or a regulator in one jurisdiction and a regulator in another, can differ in this way, and the relation locates the difference without resolving it.

The third limitation is the regress of evaluators. By the third proposition the use of a procedure does not establish its standing, and a procedure can be assessed by another procedure. It does not follow that every assessment requires an unending chain, and it does not follow that the chain can be closed by proof. In formal systems the chain ends in stipulated rules, and a result is relative to them. In institutions the chain ends in

conventions that confer authority on a procedure, and these conventions remain open to contest and revision. In empirical inquiry it ends in practices of replication and error control that are justified partly by their record and partly by the absence of better alternatives. The framework requires that the place where the chain ends be named, and does not claim that naming it settles anything. A document that said otherwise would be offering a foundation it does not possess.

The fourth limitation is that the five transition types are not claimed to be exhaustive. Repair, in which a defect in the record is corrected, revocation, in which standing that was granted is withdrawn, and delegation, in which a body assigns part of its authority to another, are plausible further types. Each interacts with the five given here and with the commitments that accumulate around them, and none is treated, since adding them would require propositions that this document does not have. The first case touches repair informally, and the treatment there should be read as an application of the existing vocabulary and not as a theory of the new type.

The fifth limitation is the relation of the results to the models in which they are proved. The propositions are exhibited by countermodel in small systems, and the programs reproduce them. A countermodel shows that a general claim of preservation, independence or recoverability cannot be made without further assumptions, and it does not show how often the failures it exhibits occur in any real system. The cases apply the vocabulary to real subjects and are as accurate as the vocabulary allows, and none of them claims that the toy models capture those subjects in detail.

The document is for a limited purpose. It makes an existing body of work intelligible as one research program, with one vocabulary and a small set of structural results, and it gives readers a way to see which parts of the corpus are instances of those results, which supply formal machinery that the results only state, and which are separate lines of inquiry. A reader who has read only this document should be able to identify the five transitions, state what standing is relative to, and recognize in a new case which of the five propositions is at issue. What each later essay contributes beyond that is the subject of the accompanying map.

References

- [1] *Lujan v. Defenders of Wildlife*, 504 U.S. 555 (1992).
- [2] F. C. Bartlett, *Remembering: A Study in Experimental and Social Psychology*, Cambridge University Press, 1932.
- [3] E. F. Loftus and J. C. Palmer, “Reconstruction of automobile destruction: An example of the interaction between language and memory,” *Journal of Verbal Learning and Verbal Behavior*, vol. 13, no. 5 (1974), pp. 585–589.
- [4] Git Project, *git-rebase: Git Documentation*, sections on recovering from an upstream rebase and rewriting published history. <https://git-scm.com/docs/git-rebase>.
- [5] Flyxion, *The Cost of Future Admissibility*, 2026. https://standardgalactic.github.io/admissibility-lab/cost_of_future_admissibility.pdf.

- [6] Flyxion, *Contrast as Control*, 2026. <https://standardgalactic.github.io/diatrube/contrast-as-control.pdf>.
- [7] Flyxion, *Semantic Geometry and Corrective Navigation*, 2026. <https://standardgalactic.github.io/memnet/corrective-navigation/semantic-geometry-spec.pdf>.
- [8] Flyxion, *Authorization and Its Remainder*, 2026. https://standardgalactic.github.io/sandbox/authorization_and_its_remainder.pdf.
- [9] Flyxion, *Echo Chamber*, 2026. <https://standardgalactic.github.io/echo-chamber/echo-chamber.pdf>.