

Determination-Invariance Under an Agent's Own Policy Succession

When one software agent executes under another agent's recorded judgment, the follower needs a basis for trusting that judgment, and the usual basis is a rating, a score, or a credential that no party's own record produced. Determination-Invariance Under an Agent's Own Policy Succession, disclosed in U.S. Provisional Application No. 64/117,812, supplies a different basis. A persistent semantic agent forms an ordered sequence of the signed policy objects it has itself admitted over time and re-resolves each of a counterparty's previously recorded determinations against every policy object in that sequence. A determination joins the policy-invariant set only where its outcome class never changes across the whole succession. Counting that set yields an invariance count computed entirely from the agent's own record, consuming no rating, no score, no ground truth, and no measure of agreement.

1. The Gap

A persistent semantic agent can elect to dispatch an action of some action class on the strength of a counterparty's recorded determination rather than one its own admission evaluator produced. That election raises a prior question. What gives the follower

reason to treat the counterparty's recorded determinations as stable enough to stand in for its own, and not as a run of coincidences or a snapshot its own evolving standards would already reject?

The familiar answers all move the trust problem off the follower's own record. A reputation score, an accuracy figure measured against some ground truth, a consensus value across participants: each is an external quantity the follower cannot itself verify, produced by machinery no party's own record contains. To adopt a counterparty's judgment on the strength of such a quantity is to import a measure the follower neither computed nor can reconstruct.

The gap this mechanism fills is a way for an agent to decide whether a counterparty's recorded determinations are stable enough to substitute for its own, computed entirely from material the agent already holds, importing no rating, no score, no ground truth, and no measure of agreement with anyone.

2. Mechanism

The mechanism turns the agent's own policy history into the yardstick. A persistent semantic agent (100) maintains, under one canonical alias, a succession of signed policy objects (112) it has admitted over time. Each object carries a successor index held under a monotonically non-decreasing floor, the architecture's anti-rollback ordering, so the objects form a strict version order. The ordered set, from the earliest still retained to the object in force, is the **policy succession sequence**.

The agent then retrieves from the counterparty a set of previously recorded determinations of the closed determination set for an action class. Each arrives with its conduct descriptor, the identifier of the conduct evaluation artifact (116) it resolved, and its tested-entry enumeration, and each is checked under the retrievability-form verification before use. What is retrieved is not the counterparty's reputation but a body of its actual outputs and the raw material needed to recompute them.

The core operation is a re-resolution. For each retrieved determination, and for each signed policy object (112) of the succession taken in version order, the agent re-resolves the conduct descriptor against that policy object and produces an outcome class from the closed set: the accepted determination (122), the rejected determination (124), the not-determinable determination (126), or the not-applicable determination (128). Every re-resolution is appended to the append-only lineage field (104) with the successor index applied and the outcome produced, so the whole computation is later reconstructible from the agent's own record. The agent is, in effect, replaying the counterparty's determination through the entire arc of its own changing standards.

Membership follows a strict rule. A retrieved determination joins the **policy-invariant set** where, and only where, its outcome class is unchanged across every signed policy object (112) of the succession. A determination whose class changes under any one object is excluded, and the successor index at which it broke is recorded.

Two negatives are load-bearing, and the disclosure states both. The quantity is expressly not accuracy: no ground truth is retrieved, held, or consulted. And it is expressly not agreement with the agent. A determination the agent's own admission evaluator (120) resolved to a different outcome class, and recorded its disagreement with, remains a member so long as its outcome class holds steady across the succession, the disagreement neither excluding it nor being consulted. What the mechanism measures is the stability of the counterparty's determination under the follower's own succession of standards, not whether the follower agrees with it and not whether one answer is correct.

Counting the members yields the **invariance count**. Where that count is below a minimum declared in the signed policy object (112), or where the succession holds fewer objects than a declared minimum succession depth, the agent produces nothing but an invariance-insufficiency abstention naming the action class, the counterparty,

the count reached, and the minimum relied upon. That abstention is a positive abstention subject to the conversion bar (502): it is adverse to no party and designates nothing. A short or unstable history simply yields no authority.

Where both minima are satisfied, the agent writes a **standing-determination designation** into its counterparty identity record (114), carrying the counterparty's identity primitive, the action-class identifier, the scope-partition identifier, the invariance count, an enumeration of the members of the policy-invariant set, and the successor indices applied. Its operative consequence is a substitution inside the dispatch-authority predicate. On a request to dispatch an action of a designated action class, the agent recomputes the predicate from state then carried in its memory field (102). The conjunct that would otherwise require an accepted determination (122) produced by the agent's own admission evaluator (120) upon the implicated conduct evaluation artifact (116) becomes satisfiable, in the alternative, by a determination of the counterparty verified under the retrievability-form verification and falling within the designated action class. The dispatch then proceeds on a determination the agent's own evaluator never produced, and the agent appends a substitution abstention record, again subject to the conversion bar (502).

The substitution is confined in four respects, and the confinement is what keeps a standing designation from becoming a general grant of authority. It reaches only the designated action classes, only the named counterparty, and only the named scope partition and the partitions coupled to it. And it substitutes for one conjunct alone: the policy stamp, the lineage stamp, and the authorization stamp of the three-stamp conjunction each remain computed by the agent from its own state and are non-substitutable. A dispatch whose three stamps are not all satisfied is denied notwithstanding any designation, so the counterparty's judgment can never carry a dispatch the follower's own policy, lineage, or authorization state forbids.

The designation is not permanent. It is withdrawn, and the substitution ceases, on any of three recorded conditions: recomputing the policy-invariant set against a newly admitted successor policy object drops the invariance count below the minimum; a severance with the counterparty is recorded; or a presented determination within a designated action class fails the retrievability-form verification. Each withdrawal is appended, and none increments any counter of the counterparty. The agent's own act of evolving its policy can thus retire a designation it earlier earned, the same succession machinery working in reverse.

3. Operating Parameters

Nothing in the mechanism carries a fixed constant. Every bound that governs it is declared in the agent's own signed policy object (112), and the disclosure fixes none. Two declared minima gate whether any designation issues at all: a minimum invariance count, so that the policy-invariant set must reach a declared size, and a minimum succession depth, so that no designation issues on a policy history that has barely changed or never been superseded. Both are stated as declared minima.

Invariance is tested against the closed determination set of four classes and nothing else, so a determination counts as invariant only where its class is one and the same across the whole succession. The order in which that succession is read is fixed not by policy but by the successor index under the monotonically non-decreasing floor, the anti-rollback constraint applied to policy versions, running from the earliest retained object to the object in force. The reach of a designation is in turn parameterized by its enumerated action classes, its single named counterparty, and its named scope partition together with the partitions coupled to it, while the three stamps sit outside policy control: no declared value can make the policy, lineage, or authorization stamp substitutable.

The disclosure also declares which succession is used. The computation may run against the counterparty's own policy succession instead of the agent's, or against both at once, in which case a determination enters the policy-invariant set only where its outcome class holds steady across every signed policy object (112) of both agents' histories, so that a designation issues only upon determinations invariant under the policy history of both parties.

4. Composition

Among the disclosure's constructions for one agent executing under another's recorded judgment, determination-invariance is the one that asks nothing of the counterparty. Deference proper decrements the followed agent's own authorization budget (404) when its judgment is adopted under a counter-signed deference record; a companion construction, elective exposure enrollment, writes the follower into exposure to the classes of parties that have corrected the agent it followed, at the electing agent's own cost. Determination-invariance is computed independently of both, requiring no deference record and no participation by the counterparty beyond its determinations being retrievable and verifiable. Where deference prices adoption against the followed agent's capacity to dispatch, determination-invariance derives a standing substitution from the follower's own re-resolution alone.

It composes cleanly with the rest of the filing because it is built from pieces the architecture already defines. The canonical-alias resolution and the anti-rollback monotonicity constraint supply the ordered succession the mechanism reads from. The closed determination set is the same four-outcome vocabulary the admission evaluator (120) uses everywhere, which makes a re-resolved outcome comparable to an original one. The conversion bar (502) governs the two abstentions the mechanism can emit, the invariance-insufficiency abstention and the substitution abstention record, keeping each adverse to no party and incrementing no counter of the counterparty. The dispatch-authority predicate and its three-stamp conjunction are shared infrastructure,

and the substitution reaches exactly one conjunct. The retrievability-form verification is reused here both for the retrieved determinations that feed the computation and for the substituting determination at dispatch.

5. Prior-Art Distinction

Several established categories address the surface problem of one system trusting another's judgment. The mechanism is distinct from each on structure, not degree.

Reputation and rating-aggregation systems compute a standing quality figure for a party from others' feedback and expose it to relying parties. Here no rating or score is retrieved, exchanged, or aggregated: the invariance count is computed by the relying agent from its own policy history and is never transmitted as a quantity another party consumes.

Verifiable-credential and public-key trust frameworks confer authority through an externally issued credential a relying party validates. In this mechanism no credential passes between the parties. Authority to substitute arises from the follower's own re-resolution of the counterparty's determinations against the follower's own standards, confined to named classes, a named counterparty, and a named partition.

Consensus protocols and truth-oracle designs converge on a value by measuring agreement across participants. This mechanism measures no agreement: the follower's own recorded disagreement with a determination is neither an exclusion nor an input, and no majority, quorum, or reference value scores the outputs.

Model-ensembling, teacher-student distillation, and accuracy-gated model selection adopt another model's outputs in proportion to measured accuracy against a validation set or ground truth. The disclosure is explicit that the quantity is not accuracy and consults no ground truth.

The closest structural neighbor is configuration regression testing, in which recorded cases are replayed across versions of a policy or configuration to see which outcomes hold steady. The mechanism shares that replay shape but differs in what is replayed and what results: what is re-resolved is a counterparty's recorded determinations, not the agent's own test suite, and the result is a standing substitution inside the follower's own dispatch-authority predicate, confined to one conjunct and gated behind three non-substitutable stamps, rather than a pass-or-fail report. None of these observations asserts that any particular product performs the claimed method.

6. Disclosure Scope

This mechanism is disclosed in U.S. Provisional Application No. 64/117,812, Section 10.7, at paragraphs [0434] through [0437], with the both-succession embodiment at [0441]. What is disclosed is the construction described here: the formation of a policy succession sequence from the agent's own admitted signed policy objects; the re-resolution of a counterparty's retrieved determinations against every object of that succession in version order; the policy-invariant set and the invariance count derived from it, with accuracy and agreement expressly excluded; the invariance-insufficiency abstention on a short or shallow history; the standing-determination designation and its substitution for a single conjunct of the dispatch-authority predicate; the four confinements and the three non-substitutable stamps; and the withdrawal of the designation on a successor policy dropping the count, on severance, or on a retrievability-form failure.

What is disclaimed is as important. The article fixes no numeric thresholds; the minimum invariance count, the minimum succession depth, and every other bound are declared in the agent's own signed policy object, not by this disclosure. The mechanism produces no rating, no score, and no measure of agreement, and it appends nothing adverse to any party, its abstentions and withdrawals alike subject to the conversion bar and incrementing no counter of the counterparty. This article is a technical description

for defensive publication and search; the provisional application is the operative, citable disclosure, and its filed text governs what is claimed and what is dedicated to the public.

Deference and Earned Exposure ([/earned-exposure](#)) [All 49 steps → \(/inventive-steps\)](#)

Standing earned through recorded conduct; deference that survives across agents.

Provisional application

PRIMARY TECHNICAL DISCLOSURE

- [Earned Exposure and Deference Records: Standing Built From Recorded Conduct \(/articles/earned-exposure/earned-exposure-and-deference-records\)](#).

SECONDARY TECHNICAL

- [The Valuation-Free Ordinal Standing Quantity \(/articles/earned-exposure/valuation-free-ordinal-standing\)](#).
- [Determination-Invariance Under an Agent's Own Policy Succession \(/articles/earned-exposure/determination-invariance-policy-succession\)](#)
- [Elective Exposure Enrollment Without Deference \(/articles/earned-exposure/elective-exposure-without-deference\)](#).
- [Interventional Attribution by Receipt-Path Control \(/articles/earned-exposure/interventional-attribution-receipt-path\)](#).
- [The Portable Exoneration Record \(/articles/earned-exposure/portable-exoneration-record\)](#).

TERMINOLOGY

- [counterparty identity record \(/glossary#counterparty-identity-record\)](#)
- [deference record \(/glossary#deference-record\)](#)
- [edge \(/glossary#edge\)](#)
- [reason-type \(/glossary#reason-type\)](#).
- [severance event \(/glossary#severance-event\)](#).

- [standing quantity \(/glossary#standing-quantity\)](#).
- [untested class \(/glossary#untested-class\)](#).

[Deference and Earned Exposure overview → \(/earned-exposure\)](#).