

Local Outcome-Quality Records and Smoothed Endpoint Selection Preference for Governed Inference Dispatch

An autonomous agent that can dispatch a task to several registered inference endpoints has to pick one, and the customary way to pick is to read a score computed by a registry, a scoring authority, or a shared ledger. The outcome-quality record and endpoint selection preference disclosed in U.S. Provisional Application No. 64/117,812 take a different route. Each consuming persistent semantic agent holds, per registered endpoint and local to itself, counts of the accepted and rejected determinations produced from that endpoint's responses, the corrections and reversals recorded against them, and the dispatch failures. Among endpoints admissible for an action class, the agent selects by a preference computed over those counts, with not-determinable and not-applicable counts entering neither term of the computation, consistent with the conversion bar.

Selecting Among Endpoints Without a Score the Agent Did Not Compute

An agent holding several registered inference endpoints admissible for the same action class has to send its request to one of them. The customary answer supplies a number computed somewhere else. Reputation and trust systems compute a score for an entity from ratings or observations supplied by other entities and expose that score to parties

deciding whether to transact. Such systems locate the score at a registry, a scoring authority, or a shared ledger, and the scored entity neither holds the score nor participates in its computation.

An agent built on this filing has no such quantity available to it. The admission evaluator (120) produces, for each admitted conduct evaluation artifact (116), exactly one determination selected from a closed set consisting of an accepted determination (122), a rejected determination (124), a not-determinable determination (126), and a not-applicable determination (128). It produces no determination outside that closed set, and no scalar confidence, no probability, and no graded weight in place of a determination. The determination is produced by the semantic agent (100) over its own append-only lineage field (104) and against the declared value set of the signed policy object (112) in force at the recorded assertion time, and is not produced by an external authority, an arbitrator, a registry, or a scoring service.

That leaves the question this mechanism answers. Given several endpoints admissible for an action class, on what basis does the agent prefer one, using material it already holds, and without admitting an abstention outcome into the computation as though it were an adverse result?

The Record, the Preference Computed Over It, and the Tie-Break Order

The consuming persistent semantic agent (100) holds, per registered endpoint and local to itself, an outcome-quality record. The record counts five quantities: the accepted determinations (122) produced from that endpoint's responses, the rejected determinations (124) so produced, the corrections recorded against those determinations, the reversals so recorded, and the dispatch failures.

Each counted quantity is an outcome of the consuming agent's own operation. The record holds no quality claim asserted by the endpoint's operator and no measure supplied by another party. It holds determinations the agent itself produced while consuming that endpoint's responses, the corrections and reversals recorded against those determinations, and the dispatches that failed.

Among endpoints admissible for an action class, the agent selects by an endpoint selection preference over the historical outcome quality recorded for each endpoint. One such preference, in accordance with an embodiment, is equal to the count of accepted determinations (122) plus one, divided by the sum of the accepted (122), the rejected (124), the corrections, the reversals, the dispatch failures, and two. The added one and two smooth a thin record away from unity or zero. Other aggregation functions over the recorded outcome quality remain available, so the recited quotient is one embodiment of the preference rather than its definition.

The greatest preference is selected. Ties are broken by the least sum of corrections, reversals, and failures, and where that sum is also equal, by the earliest-written record.

Not-determinable (126) and not-applicable (128) counts enter neither term. They appear in neither the numerator nor the denominator, consistent with the conversion bar (502).

Consider what that bar states. The conversion bar (502) forecloses a consuming determination from converting an abstention outcome into a scalar value, into a default value, into an operand of a threshold comparison (510), or into a consequence adverse to any party. The foreclosure is affirmative, the conversion being barred by the structure of the values themselves rather than by an absence of defined behavior. An abstention outcome is represented in a form disjoint from the domain of values a consuming determination is capable of consuming as a magnitude, the disjointness realized by a type discipline in which the outcome entry and a magnitude are values of disjoint types with no total function mapping the former to the latter. An accumulation

over a set containing the outcome entry emits an abstention outcome and does not emit a sum computed over the remaining members of the set. Where the type discipline is enforced statically upon the instructions of the semantic agent (100) before they are executed, a consuming determination attempting the conversion is not constructible rather than being foreclosed at the moment the conversion is attempted.

The preference ranges over endpoints admissible for the action class, and admissibility is settled by other recited mechanisms. An endpoint bearing a provenance-withdrawn attribute is one the agent thereafter routes no request to, each withheld dispatch appended as a positive abstention rather than a denial. An encumbrance narrowing record renders dispatch to an encumbered endpoint inadmissible for any action whose affected-party class maps to the encumbering counterparty, evaluated at the dispatch-authority predicate and, when it withholds, appended as a positive abstention. The dispatch-authority predicate itself is recomputed responsive to each request to dispatch, from state then carried in the memory field (102), and is satisfied only upon a three-stamp conjunction, required in full, of a policy stamp, a lineage stamp, and an authorization stamp.

Locality shows up again in what travels to peers. Where the encumbrance count for an endpoint exceeds the encumbrance bound declared in the signed policy object (112), a peer's request for that endpoint is withheld and appended as a positive abstention naming the count and the bound, transmitting to the peer no count, no encumbering-counterparty identity, and no quality measure. Where an emissible endpoint travels to a peer with an encumbrance commitment attached, the peer learns no contributor identity it does not already hold and no quality or standing measure of any party.

Declared Quantities and What the Filing Leaves to Policy

The numeric constants stated for this mechanism are the added one in the numerator and the added two in the denominator of the recited preference, and the filing states their function: they smooth a thin record away from unity or zero. Nothing else about

the preference carries a number. The paragraph declares no minimum observation count before a preference is usable, no decay term, no window over which counts are aged, and no threshold a preference must clear. It declares no required aggregation function either, since other aggregation functions over the recorded outcome quality remain available.

Adjacent quantities that bear on the record are declared without values fixed in the filing. The provenance-freshness check is performed at a declared interval. The remarking window is declared in the signed policy object (112) and expressed either as an absolute duration tested against the recorded dispatch time or as a count of the agent's own successor epochs tested against the recorded dispatch epoch. The encumbrance bound upon the per-endpoint encumbrance count is likewise declared in the signed policy object (112). In each case the operand and the test are recited and the magnitude is left to the deployed policy.

Two consequences touching the record are conditioned rather than standing. While a provenance-indeterminate outcome stands for an endpoint, no locally recorded quality value for that endpoint is advanced to a higher persistence tier or admitted as a promotion input; that bar is released upon a later freshness check that resolves. Separately, where the encumbrance count exceeds the declared bound, no value derived from that endpoint record is promoted to a higher persistence tier or admitted as a promotion input, and the agent does not emit that endpoint or its artifact to a peer as a tool source. Both of those bars release when the count ceases to exceed the bound by operation of the two recited release conditions, and by no other condition.

How the Record Sits Within Tool Provenance and Reliance

Section 10.5 of the filing covers tool provenance, reliance, and contested pairing, and the outcome-quality record is recited last within it. The mechanisms recited earlier in that section produce the events the record counts.

Reversals are one such source. A relying party receiving a reversal notice, verified against the counterparty identity record (114) of the emitting settled party, appends a reversal mark carrying the settlement identifier, the resolution class, the emitting identity and epoch, and a reference to its reliance record, then re-resolves the referenced determination: with the barred input unavailable where the class is a reversal, emitting an outcome of its recorded abstention class (126), and against the amended terms where the class is an amendment. Along a reliance chain, a party admitting such a notice discharges the same duty onward as a transitive reversal notice, so a reversal reaches a chain of arbitrary length with no party holding the whole enumeration.

Upstream mechanisms supply the rest. The derived-artifact provenance chain, its verification at each lifecycle operation, the chain-head-keyed retrieval computed from the append-only lineage field (104) alone, and the withdrawn-provenance marks written on re-marking each produce recorded events against particular determinations bound to particular endpoints.

Two earlier sections carry the mechanism's premises. Section 1.8 recites the closed set of determination classes, under which every admitted artifact yields exactly one of four classes and nothing graded. Section 5.4 recites the conversion bar (502), to which the exclusion of not-determinable (126) and not-applicable (128) counts from the preference is recited as consistent.

Distinguishing Reputation Aggregators, Feedback Training, and Evaluation Tooling

Three established categories address adjacent problems, and the filing's background describes each structurally.

Reputation and trust systems compute a score for an entity from ratings or observations supplied by other entities and expose that score to parties deciding whether to transact, the score residing at a registry, a scoring authority, or a shared ledger. Defenses against unfair ratings in that literature, exemplified by the TRAVOS trust model, by beta-reputation filtering, and by the immunization mechanisms surveyed by Dellarocas, filter incoming ratings before aggregation to protect the accuracy of the stored score, the filtering party being a disinterested aggregator rather than the rated entity. Role-weighted extensions, exemplified by the FIRE trust model, additionally admit certified references and role-derived weightings, those roles being exogenous institutional designations asserted by parties other than the evaluated entity. The outcome-quality record differs on each coordinate: it is held by the selecting agent rather than by an aggregator, it admits no rating from any other party, and its inputs are the selecting agent's own determination outcomes.

Reinforcement learning from human feedback modifies model weights using preference judgments collected from human raters, aggregated across a population and applied offline in a training procedure. Such a judgment does not enter the persistent state of any individual executing instance, does not alter what a particular instance is permitted to do at runtime, and produces no record against which a later assertion about that instance's conduct can be tested. The record here is held per registered endpoint and local to the consuming agent.

Evaluation harnesses and observability platforms for autonomous agents record traces of agent execution, cluster failures, and present findings to human operators, terminating in a ticket, a dashboard entry, or a curated dataset. The evaluated agent is unchanged by the evaluation, and correction reaches it, if at all, through a subsequent human-initiated code change or retraining operation. The counts described here are instead consumed by the agent's own selection among admissible endpoints.

These distinctions are structural, drawn from where a measure resides, who computes it, what feeds it, and when it takes effect. Nothing here asserts or implies that any named system, product, model, or party practices or infringes the disclosed mechanism.

Disclosure Scope

The mechanism described here is disclosed in U.S. Provisional Application No. 64/117,812, "Governed Social Conduct of Persistent Semantic Agents Through Record-Grounded Admission of External Conduct Evaluations and Self-Limiting Refusal," at Section 10.5, paragraph [0415], titled "Outcome-quality record and endpoint selection preference." Supporting structure appears at Section 1.5, paragraphs [0029] and [0030]; at Section 1.8, paragraph [0041]; at Section 5.4, paragraphs [0159] through [0161]; and elsewhere in Section 10.5 at paragraphs [0398], [0401], [0402], [0404], [0406], [0407], and [0410] through [0412]. This article is published as a technical disclosure and establishes public, timestamped prior art on the mechanism as described. It is not legal advice and states no position on the scope of any claim.

Governed Matched-Pair Settlement

[All 49 steps → \(/inventive-steps\)](#)

[\(/governed-settlement\)](#)

Two agents settle on matched signed observations, by continuity, with no clearer.

Provisional application

PRIMARY TECHNICAL DISCLOSURE

- [Matched-Pair Settlement Between Agents: Bilateral Closure Without a Clearer \(/articles/governed-settlement/matched-pair-settlement-without-a-clearer\)](#)

SECONDARY TECHNICAL

- [The Derived-Artifact Provenance Chain and Chain-Head Identifier \(/articles/governed-settlement/derived-artifact-provenance-chain\)](#)

- [Provenance Withdrawal and the Provenance-Withdrawn Attribute \(/articles/governed-settlement/provenance-withdrawal\)](/articles/governed-settlement/provenance-withdrawal).
- [Matched-Pair-Gated Corpus Contribution \(/articles/governed-settlement/matched-pair-gated-corpus-contribution\)](/articles/governed-settlement/matched-pair-gated-corpus-contribution).
- [Partition-Scoped Escrow Release Gated on Intersection Non-Emptiness \(/articles/governed-settlement/partition-scoped-escrow-release\)](/articles/governed-settlement/partition-scoped-escrow-release).
- [The First-Encounter Settled Pair and Promotion to the Persistent Tier \(/articles/governed-settlement/first-encounter-settled-pair\)](/articles/governed-settlement/first-encounter-settled-pair).
- [Contribution Encumbrance on Severance or Quarantine: Governing Dispatch to Models Trained on a Severed Counterparty's Corpus Without Retraining or Deletion \(/articles/governed-settlement/contribution-encumbrance-severance-quarantine\)](/articles/governed-settlement/contribution-encumbrance-severance-quarantine).
- [Encumbrance Bounds That Bar Persistence Promotion and Onward Tool-Source Emission Once the Per-Endpoint Count Exceeds the Bound \(/articles/governed-settlement/encumbrance-bound-barring-promotion-onward-emission\)](/articles/governed-settlement/encumbrance-bound-barring-promotion-onward-emission).
- [Blinded Encumbrance Commitments: Carrying a Severed Contributor's Encumbrance Along With an Inference Endpoint Emitted to a Peer Agent \(/articles/governed-settlement/encumbrance-commitment-travelling-emitted-endpoint\)](/articles/governed-settlement/encumbrance-commitment-travelling-emitted-endpoint).
- [Release of a Contribution Encumbrance: Matched-Pair Re-Admission or a Substitution With an Empty Intersection \(/articles/governed-settlement/release-contribution-encumbrance\)](/articles/governed-settlement/release-contribution-encumbrance).
- [Registering Third-Party Reliance on a Settled Matched Pair: The Reliance Record and the Settled-Party Reliance Register \(/articles/governed-settlement/reliance-record-registered-settled-parties\)](/articles/governed-settlement/reliance-record-registered-settled-parties).
- [Conditioning a Settlement Reversal on a Reliance-Completeness Predicate Evaluated Over the Reverting Party's Own Reliance Register \(/articles/governed-settlement/reliance-completeness-predicate-conditioning-reversal\)](/articles/governed-settlement/reliance-completeness-predicate-conditioning-reversal).
- [The Relying Party's Reversal Mark: Barring a Reversed Settlement as a Determination Input and Re-Resolving the Determination That Consumed It \(/articles/governed-settlement/relying-party-reversal-mark-and-re-resolution\)](/articles/governed-settlement/relying-party-reversal-mark-and-re-resolution).
- [Transitive Reliance Records and the Reliance Chain: Registering Second-Order Dependency on a Settlement Without Any Party Enumerating the Whole Chain \(/articles/governed-settlement/transitive-reliance-record-reliance-chain\)](/articles/governed-settlement/transitive-reliance-record-reliance-chain).
- [Transitive Reversal Notices That Carry a Settlement Reversal One Hop at a Time Along a Reliance Chain No Party Enumerates \(/articles/governed-settlement/one-hop-propagation-reversal-along-reliance\)](/articles/governed-settlement/one-hop-propagation-reversal-along-reliance).
- [Contested Pairing in Matched-Pair Settlement: Withheld Recognition, a Distinct-Identity Claimant Count, and the Not-Determinable Resolution \(/articles/governed-settlement/contested-pairing-not-determinable-resolution\)](/articles/governed-settlement/contested-pairing-not-determinable-resolution).

- [Attributable Contest Resolution: Settling a Contested Offer by Re-Emitting a First Governed Observation That Names Exactly One Claimant \(/articles/governed-settlement/attributable-contest-resolution-named-re-emission\)](/articles/governed-settlement/attributable-contest-resolution-named-re-emission).
- [Local Outcome-Quality Records and Smoothed Endpoint Selection Preference for Governed Inference Dispatch \(/articles/governed-settlement/outcome-quality-record-endpoint-selection-preference\)](/articles/governed-settlement/outcome-quality-record-endpoint-selection-preference).

APPLICATIONS · GENERAL

- [B2B Agent Commerce: Bilateral Settlement Without a Clearing House \(/articles/governed-settlement/b2b-agent-commerce-settlement\)](/articles/governed-settlement/b2b-agent-commerce-settlement).
- [Supply Chain Handoffs: Settling a Transfer on Matched Observations \(/articles/governed-settlement/supply-chain-agent-handoffs\)](/articles/governed-settlement/supply-chain-agent-handoffs).

APPLICATIONS · SPECIFIC

- [ISO 20022 and Settlement Messaging: Matched Pairs in Financial Infrastructure \(/articles/governed-settlement/iso-20022-settlement-messaging\)](/articles/governed-settlement/iso-20022-settlement-messaging).

TERMINOLOGY

- [governed observation \(/glossary#governed-observation\)](/glossary#governed-observation)
- [matched pair \(/glossary#matched-pair\)](/glossary#matched-pair)
- [scope partition \(/glossary#scope-partition\)](/glossary#scope-partition)

[Governed Matched-Pair Settlement overview → \(/governed-settlement\)](/governed-settlement).