

Interventional Attribution by Receipt-Path Control

A confirmation latency recorded against a counterparty carries two explanations the recorded number cannot separate: the counterparty is slow, or the domain it operates in is slow and every party there confirms slowly. An agent that attributes the delay from the latency alone records it against a counterparty whose only recorded property is a slow corridor. The interventional attribution mechanism, disclosed in U.S. Provisional Application No. 64/117,812, resolves the ambiguity by varying a property of the semantic agent's own dispatch rather than by watching passively. It routes matched interactions to close either through a receipt the counterparty controls or through a receipt it does not, assigns the two paths by a rule committed before any outcome is seen, and reads the latency difference between them to attribute the delay to the counterparty or to the domain.

1. The Gap

A confirmation latency recorded against a counterparty admits two explanations that are not distinguishable from the recorded latency alone. In the first, the counterparty is slow to confirm. In the second, the domain in which the interaction occurs is slow, and every party operating within that domain confirms slowly. The two produce the same number.

An agent that reads the number and stops has to guess. Attribute the delay to the counterparty, and it writes an adverse mark against a party whose only recorded property is that it operates in a slow corridor. Assume the domain instead, and it lets a genuinely unresponsive counterparty hide behind the corridor. Neither guess is grounded, because the two causes are indistinguishable in the one signal the agent holds.

Separating the two by watching passively would mean comparing the counterparty against some outside measure of how its domain performs, and that measure is not in the agent's own records. Importing it would mean trusting figures the counterparty or a third party supplies, which the agent cannot verify. So the passive route stays closed: the recorded latency does not carry the distinction, and the agent will consult no counterparty-supplied score to recover it.

The mechanism of paragraph [0358] of U.S. Provisional Application No. 64/117,812 takes the other route. Instead of adding data, it adds an intervention. The semantic agent (100) varies a property of its own dispatch and reads how the recorded latency responds. Because the agent controls the variable, it does not need anyone else's records, and because the variation is committed in advance, the result is verifiable after the fact.

2. Mechanism

The variable the agent controls is the receipt that closes an interaction. Every dispatched interaction closes upon some receipt, and receipts fall into two kinds by who controls when they are produced.

A **first closure path** closes upon a receipt the counterparty controls: an attestation or a signed confirmation the counterparty itself produces. The counterparty decides when that receipt exists, so any slowness in producing it is the counterparty's own, layered on top of channel overhead that has nothing to do with counterparty conduct.

A **second closure path** closes upon a receipt the counterparty does not control: an execution outcome recorded by the semantic agent (100) upon completing its own dispatched operation, or an acceptance recorded by the principal to which the agent is bound. The counterparty does not determine when this receipt is produced, so its latency reflects the domain rather than the counterparty's own responsiveness in confirming.

The agent varies only a property of a dispatch it is already making. The assignment is meant to leave what the counterparty receives unchanged, changing only which receipt is treated as closing the interaction on the agent's own record. That is what makes the move an intervention on the agent's own books rather than a manipulation of the counterparty.

To compare the two paths meaningfully, the agent groups interactions into a **matched class**: interactions identical in action class, stakes, urgency, and scope partition. Within a matched class, the interactions differ in nothing the agent declares relevant except the closure path assigned to them, so a latency difference between the paths is not confounded by a difference in what was being done or how much rode on it.

Assignment within a matched class follows one of two rules. The first is **deterministic alternation**. The second is a **pseudorandom selection seeded from a value committed to the append-only lineage field (104) before the first assignment of the matched class**. Either rule is meant to leave the assignment indistinguishable to the counterparty on the dispatch it receives, and the counterparty is not notified that an assignment has occurred. Committing the seed first is the load-bearing detail: the assignment sequence is fixed before any latency is observed, so the agent cannot have chosen the sequence to produce a result it wanted after seeing the outcomes. Anyone holding the lineage can replay the seed and confirm the assignments were not back-fitted.

With the two paths populated, the reading follows. The agent computes a first closure latency and a second closure latency, each as the arithmetic mean of the dispatch-to-closure intervals for the interactions assigned to that path (paragraph [0359]). The first path carries channel overhead that has nothing to do with counterparty conduct. The agent removes that overhead by a **channel baseline**, computed as the pooled arithmetic mean of first-path closure durations recorded with counterparties that pass a qualifying predicate evaluated on the agent's own records alone: persistent-tier identity in the counterparty identity record (114) and a declared minimum count of first-path closures, each within a declared responsiveness bound. The predicate consults no counterparty-supplied score. Normalization is subtraction, bounded below at zero. The **latency differential** is the normalized first latency less the second latency.

The differential is then run through an ordered branch (paragraph [0360]). If either path holds fewer than a declared minimum per-path volume, the agent appends a not-determinable determination (126) and attributes nothing, because a handful of closures on one path cannot support a stable mean. If the differential exceeds a declared threshold, the latency is attributed to the counterparty by an attribute written to the counterparty identity record (114). If the differential does not exceed the threshold but the second latency is itself elevated above a **second-path baseline** by a declared margin, the latency is attributed to the domain and nothing adverse is written to the counterparty. The second-path baseline is the mean second latency across the agent's other matched classes, with the counterparty under evaluation excluded so the comparator is not contaminated by that counterparty's own closures. If neither condition holds, a no-adverse-signal outcome is appended.

Two guards keep the attribution honest. The threshold is bounded below by the standard deviation of the qualifying-closure durations from which the channel baseline was computed (paragraph [0361]); a threshold under that floor would attribute ordinary channel noise to the counterparty, so the decision line is tied to the measured dispersion of the channel it is read against. And a distinguishability check (paragraph

[0362]) computes counterparty conduct other than latency, the rate of non-closure, of returned exceptions, and of amendment requests, per closure path. If that conduct differs across the paths by more than a declared manipulation margin, then the assignment the agent meant to be indistinguishable was in fact distinguishable to the counterparty, the latency evidence of that class is invalidated, and the agent appends a manipulation attribute instead of attributing or exonerating the latency.

3. Operating Parameters

Every quantity below is declared in the signed policy object (112). The filed disclosure fixes the structure and the relationships, not the numbers.

- **Matched-class feature tuple.** Interactions are matched on four recorded features held identical across the class: action class, stakes, urgency, and scope partition.
- **Assignment rule.** Either deterministic alternation, or pseudorandom selection seeded from a value committed to the append-only lineage field (104) before the first assignment, so the sequence can be replayed from the recorded seed.
- **Central-tendency measure.** The arithmetic mean of dispatch-to-closure intervals, applied identically to both closure paths.
- **Channel baseline and its qualifying predicate.** A pooled arithmetic mean over first-path closures of qualifying counterparties, where qualification requires persistent-tier identity and a declared minimum count of first-path closures, each within a declared responsiveness bound. Normalization is subtraction, bounded at zero.
- **Minimum per-path volume.** A declared count below which either path yields a not-determinable determination (126) rather than an attribution.
- **Attribution threshold.** Declared, and bounded below by the standard deviation of the qualifying-closure durations, expressed in the units of the baseline and threshold.

- **Second-path baseline and elevation margin.** The baseline is the mean second latency across the agent's other matched classes, excluding the counterparty under evaluation; the margin by which the second latency must exceed it to support a domain attribution is declared.
- **Manipulation margin.** A declared bound on the permissible cross-path difference in non-latency conduct, above which the class's latency evidence is discarded.

4. Composition

The receipt-path control is the input stage of the interventional-attribution section; the differential, the channel baseline, and the ordered branch are its consumers. All of it runs on substrate the rest of the filing already defines. The signed policy object (112) supplies every parameter, so how aggressive the assignment and its thresholds are is governed by the same signed object that governs the agent's other conduct. The append-only lineage field (104) holds the committed seed and the appended outcomes, which is what makes the sequence replayable and the result verifiable. Attributions land as attributes on the counterparty identity record (114), the same record the agent's other mechanisms read and write. And when evidence is thin the branch produces the not-determinable determination (126), the filing's shared primitive for an under-evidenced outcome, rather than defaulting to something adverse.

The tie to earned exposure is the signal itself. The valuation-free ordinal standing quantity of paragraph [0336] derives one of its three timing statistics from latency, computed as the median reciprocation interval between the parties over the agent's own lineage entries. A counterparty stuck in a slow corridor produces a long reciprocation interval through no fault of its own, and a standing quantity read off raw latency would fold that domain-caused delay into the counterparty's own timing signal. The interventional attribution disciplines exactly that class of timing signal before it hardens into a judgment: latency the mechanism assigns to the domain is not latency the counterparty should carry. The mechanism does not compute a standing multiplier itself; it operates on the same recorded quantity the standing machinery consumes.

Because the whole procedure runs on the agent's own dispatch and its own records, it composes without any coordinating party, registry, or shared clock, the same self-contained posture the rest of the filing takes toward verification.

5. Prior-Art Distinction

Several established categories touch confirmation latency or randomized measurement, and the mechanism differs from each structurally rather than by degree.

Passive latency and service-level monitoring records durations and compares them against a threshold or a rolling baseline. It is observational: it reads the latency that occurred and cannot separate a slow counterparty from a slow domain, because both produce the same recorded duration. The mechanism here does not add a better threshold; it adds a controlled variation in the agent's own dispatch, which is the only thing that breaks the degeneracy the observational approach inherits.

A/B testing, randomized experiments, and bandit assignment vary a treatment presented to a subject and measure the subject's response, typically to optimize a metric. Here the varied property is not something shown to the counterparty. What varies is which receipt closes an interaction the agent is making regardless, recorded on the agent's own books; the counterparty is not notified; and the assignment is committed to an append-only log before outcomes are observed, so the sequence is auditable rather than merely random.

Causal inference and instrumental-variable estimation recover causal effects from observational or designed data. The mechanism's instrument is the agent's own choice of closure path, chosen by a pre-committed seed rather than found in the data, and the output is a bounded attribution decision with an explicit not-determinable branch, not a point estimate of an effect size.

Reputation and trust systems score parties from accumulated outcomes. This mechanism sits upstream of any such score: it guards the latency input by attributing it correctly first, declines to write anything adverse when volume is insufficient or when the distinguishability check shows the counterparty could tell the paths apart, and refuses to attribute channel noise by tying its threshold to the measured dispersion of the channel.

None of the above turns on any named product or party, and nothing here implies that any such system infringes; the distinctions are of mechanism, not of market.

6. Disclosure Scope

This article describes subject matter disclosed in U.S. Provisional Application No. 64/117,812, Section 10.2, principally paragraph [0358], "Interventional attribution by receipt-path control," together with the supporting paragraphs on which its attribution depends: the latency differential and channel-baseline normalization of [0359], the attribution-versus-exoneration branch of [0360], the standard-deviation lower bound on the threshold of [0361], and the distinguishability check of [0362]. What is disclosed is the structure: closure through a counterparty-controlled or non-counterparty-controlled receipt, matching on action class, stakes, urgency, and scope partition, assignment by deterministic alternation or by a pre-committed seeded selection, and the reading of a normalized latency differential against policy-declared thresholds to attribute latency to a counterparty or to a domain.

What is not disclosed here, and is not to be inferred from this article, includes every specific numeric parameter value; each parameter named above is declared in the signed policy object (112) and no value is fixed by this publication. The portable, transmissible form of an exoneration record is the subject of a separate disclosure and is not treated here. Any further sub-mechanism, deployment, or variant not set out in the filed provisional is outside the scope of this article, and nothing in it should be read to place such material in the public domain.

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 \(/articles/earned-exposure/termination-invariance-policy-succession\)](#)
- [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).