

Trustpilot Signals and Conduct-Derived Admissibility

A software agent arrives at a counterparty it has never contacted, carrying a number computed elsewhere by rules the receiving side cannot inspect. Consumer review platforms such as Trustpilot solve the human version of this problem well: they collect reviews from people who dealt with a business and publish them, with a summary rating, where a prospective customer will see them. The earned-exposure mechanism of U.S. Provisional Application No. 64/117,812 addresses a different question in the same category. It moves no rating, score, or standing quantity between agents at all. Instead, counter-signing another agent's determination decrements the followed agent's own authorization budget, so the cost of being followed rises with adoption, and what the deferring agent earns is admissibility rather than reputation.

The Score Travels, the Conduct Does Not

A software agent contacts a counterparty it has never dealt with, holding a number computed elsewhere, over encounters the receiving side cannot inspect, under rules it did not choose. That side must now decide whether the number stands in for experience it does not have.

This is the structural weakness in any portable standing quantity, whether a star rating, a vouch, a stake, or a weighted endorsement. Detached from the encounters that produced it, the number carries an aggregation policy the receiving party cannot audit, and becomes an operand to be thresholded and averaged.

A second consequence draws less attention. In conventional delegation and capability-attenuation designs, the cost falls on the relying party and the party relied upon is protected, attenuation narrowing only the capability passed downward. Adoption is treated as costless to the adopted.

For agents that dispatch consequential actions the question is narrower than reputation: whether an evaluative artifact from a party with no prior encounter history should be admitted for consideration at all. That is a question about admissibility, and the disclosed architecture requires it to be settled on records the receiving agent holds and verifies for itself.

What Trustpilot Does, and the Category It Serves

As publicly described, Trustpilot operates an open review platform. People who have dealt with a business can write a review of that experience, the review appears on a company profile page, and a summary star rating is presented alongside. Its public materials also describe business-side capabilities, among them claiming a profile, inviting customers to leave reviews, and responding publicly, together with published guidelines and moderation addressed to fake and incentivized reviews. How any of that is implemented is not characterized here.

The design fits the need it serves. A prospective customer has no encounter history with a merchant and needs a compressed signal quickly, and aggregated first-person accounts from people who did have the encounter supply one. Publication gives the merchant a standing audience for how it behaves.

Review platforms of this kind are addressed to a human reader exercising judgment. A person reads the rating, reads several reviews, notices that complaints cluster around shipping rather than product quality, weighs the business response, and forms a view. The summary number opens that judgment rather than standing in for it. What follows is not an assessment of Trustpilot. It sets out what the disclosed architecture requires when the party weighing an evaluative artifact is software, not a person.

Counter-Signature, the Merit-Independent Decrement, and Earned Admissibility

The mechanism disclosed in U.S. Provisional Application No. 64/117,812 lets one semantic agent (100) execute under the recorded evaluative determination of a second semantic agent (100) with no rating, score, standing quantity, or credential passing between them. Its precondition is restrictive. The deferring agent already finds its dispatch-authority predicate satisfied on recomputation from its memory field (102): authorized on its own record to determine the matter for itself, it elects instead to execute under a determination held in the second agent's lineage field (104). An agent whose predicate fails, whose authorization gate (300) stands in the withheld state (310), or which has entered the non-executing cognitive mode (302) does not defer. Its holding is a withholding, not a following.

Two tests run before any effect is given. The origin-disjointness test asks whether the two agents are assignable to a common origin-equivalence class (200); where they are, the record is appended as an attempted deference record and given no effect, decrementing nothing and appending nothing adverse to the followed agent. The retrievability-form verification then checks three ordered conjuncts: continuity of the followed agent's identity against the continuity history store in the counterparty identity record (114), the tested-entry enumeration against a lineage commitment, and membership of the determination in the closed set with a matching action class. The verification is of form and retrievability, not of merits: the deferring agent does not retrieve the enumerated entries, evaluate their sufficiency, or re-derive the

determination. On satisfaction of both, a deference record (900) is appended, carrying a counter-signature window field (912) from which the followed agent obtains the applicable window.

The inversion sits in what happens next. Electing to counter-sign, the followed agent appends a deference counter-signature (902) and by that append decrements its own authorization budget (404), the budget gating that agent's own non-advisory dispatch. The decrement schedule is declared in that agent's signed policy object (112), the amount required to be not less than that declared for issuance of a conduct evaluation artifact (116) by that agent, and to be a monotonically non-decreasing function of the count of counter-signatures appended within a declared window, bounded above by a declared maximum, so the marginal cost of being followed rises with the rate of adoption. Base, increment, and bound are policy-declared, the disclosure giving step and linear forms, not fixed values.

The decrement occurs without regard to whether the deference was well founded: no adjudication of the correctness of the adopted determination, of the reasonableness of the reliance, or of the merits of the underlying artifact is performed as a condition of it. Responsive to the authorization budget (404) satisfying a floor declared in the followed agent's policy object, its authorization gate (300) is written to the withheld state (310) for an enumerated set of action classes and an escalation record is emitted to its principal. That budget is replenished by the ordinary replenishment procedure and by no other path, so being followed spends a replenishable quantity and does not terminate the agent.

What the deferring agent gets in return is not standing. Responsive to receipt of the counter-signature, and not otherwise, it retrieves a corrector-class enumeration (904): the origin-equivalence classes (200) recorded both as having corrected the followed agent in that action class and as having had that correction thereafter conformed to. A two-part conformance filter (906) admits a class only where both parts hold, so a class that corrected the agent without subsequent recorded conformity is excluded. Each

class is carried by a class-defining reference rather than by the followed agent's local identifier. The deferring agent appends an exposure edge record (910) and writes each class into its own replenishment register (402) as an admissible-but-unexercised entry.

That entry has two effects and no third. An arriving conduct evaluation artifact (116) from a party assignable to the class is admissible to the deferring agent's admission evaluator (120) notwithstanding no prior encounter history with that party, and is not held in the pre-settlement inert state on that ground. The entry is also not counted in renewal while its exercise field stands unexercised. On receipt of such an artifact the agent appends an exercised transition record (908), binding the correction received to the deference that earned the exposure through which it arrived.

Where the Two Designs Diverge

The convergence is genuine: both answer the problem of acting toward a counterparty you have no history with. The divergence is in what crosses the boundary. A review platform produces a portable summary a third party consumes, and the value is precisely that the signal travels. The filed architecture forbids that transfer. Nothing transmitted describes the followed agent's standing, and no rating, weight, score, endorsement, vouch, or standing quantity of that agent is retrieved, admitted, or consumed in computing the enumeration. The conversion bar (502) extends this: the deference abstention entry is not converted into a scalar value, a default value, or an operand of a threshold comparison, and that the deferring agent did not determine the matter for itself is not treated as a default judgment against it.

Cost runs in a different direction as well. The disclosed architecture requires the metered consequence to fall on the followed agent, in the units of its own capacity to dispatch, on a schedule non-decreasing in the count of counter-signatures appended within a declared window.

What is earned differs in kind. An admissible-but-unexercised entry is required to assert nothing about anyone: it states only that an artifact from a particular class may be evaluated, and the evaluation itself stays with the receiving agent's own admission evaluator (120). The filing further requires that the deferring agent produce no determination concerning any class it enumerates.

Coexistence, and What the Disclosure Does Not Solve

These are not competing products, and a deployment can hold both. Human-facing commercial trust is the review platform's domain, and nothing in the filed architecture replaces it. The disclosed mechanism operates between software agents, on records each agent holds and verifies for itself, gating whether an evaluative artifact enters an admission evaluator (120). A marketplace operator could present public reviews to its human users and run conduct-derived admissibility beneath them.

The disclosure also has limits, several stated in the filing itself.

It does not determine whether a determination is correct. An agent can defer to one that later proves wrong; the architecture's response is to bind the correction it subsequently receives to the deference in its own record, not to penalize anyone. Where the followed agent supersedes or retracts the adopted determination, the deference record (900) satisfies no further dispatch-authority predicate, though a dispatch already performed is not undone and is converted into a fault of neither agent.

It does not produce a global view. Assignability is evaluated on the deferring agent's own records and the arriving artifact, without a centralized registry and without coordination with a further execution node.

Nor does it resolve disagreement. Where two followed agents present determinations of differing outcome class on one conduct descriptor, the deferring agent appends a deference conflict record naming both, adopts neither, and either determines the

matter itself or withholds the dispatch as a positive abstention. It also does not reach the consumer case: a person reading about a merchant is not an agent recomputing a dispatch-authority predicate.

Disclosure Scope

The earned-exposure and deference mechanism described here, in which a deference record (900) transmitted for counter-signature decrements the followed agent's own authorization budget (404) without regard to the merits, the deferring agent enrolls that agent's conformance-filtered corrector classes into its own replenishment register (402) as admissible-but-unexercised entries, and no rating, score, standing quantity, or credential passes between them, is disclosed in U.S. Provisional Application No. 64/117,812. Amounts, windows, floors, and bounds referred to here are declared in the relevant agent's signed policy object (112); where the filing declares a quantity without a value, none is stated here. Conditioned outcomes, including the write of the authorization gate (300) to the withheld state (310), occur only where the declared condition is satisfied. The description of Trustpilot is external context drawn from public materials, not a claim of the filing. References to Trustpilot are to public materials and are used for comparison only; no relationship, endorsement, or infringement is asserted.

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

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 \(/articles/earned-exposure/earned-exposure-and-deference-records\)](/articles/earned-exposure/earned-exposure-and-deference-records)

SECONDARY TECHNICAL

- [The Valuation-Free Ordinal Standing Quantity \(/articles/earned-exposure/valuation-free-ordinal-standing\)](/articles/earned-exposure/valuation-free-ordinal-standing)
- [Determination-Invariance Under an Agent's Own Policy Succession \(/articles/earned-exposure/determination-invariance-policy-succession\)](/articles/earned-exposure/determination-invariance-policy-succession)
- [Elective Exposure Enrollment Without Deference \(/articles/earned-exposure/elective-exposure-without-deference\)](/articles/earned-exposure/elective-exposure-without-deference)
- [Interventional Attribution by Receipt-Path Control \(/articles/earned-exposure/interventional-attribution-receipt-path\)](/articles/earned-exposure/interventional-attribution-receipt-path)
- [The Portable Exoneration Record \(/articles/earned-exposure/portable-exoneration-record\)](/articles/earned-exposure/portable-exoneration-record)
- [Attribution, Exoneration, and the Domain Baseline: An Ordered Branch Over the Latency Differential \(/articles/earned-exposure/attribution-versus-exoneration-branch-domain-baseline\)](/articles/earned-exposure/attribution-versus-exoneration-branch-domain-baseline)
- [Bounded Count of Dispatches Per Deference Record: Capping How Many Executions One Counter-Signed Adopted Determination Authorizes \(/articles/earned-exposure/bounded-count-dispatches-deference-record\)](/articles/earned-exposure/bounded-count-dispatches-deference-record)
- [Coarse Lineage-Segment Granularity: Verifying a Tested-Entry Enumeration Without Disclosing Which Lineage Entries Were Tested \(/articles/earned-exposure/coarse-lineage-segment-granularity-tested-entry\)](/articles/earned-exposure/coarse-lineage-segment-granularity-tested-entry)
- [Correction-Exposure Decay: An Authorization Quantity That Depletes With Agent Operation and Renews Only on Recorded External Evaluation \(/articles/earned-exposure/correction-exposure-decay-authorization-quantity\)](/articles/earned-exposure/correction-exposure-decay-authorization-quantity)
- [Cause-Keyed Gate Restoration: A Distinct Return Path for Decay-Induced Withholding of Agent Authorization \(/articles/earned-exposure/distinct-return-path-decay-induced-withholding\)](/articles/earned-exposure/distinct-return-path-decay-induced-withholding)
- [Edge Severance Events and the Three-Valued Survival State: Resolving Whether a Counterparty Relationship Outlived a Payment-Class or Obligation-Class Termination \(/articles/earned-exposure/edge-severance-events-three-valued-survival\)](/articles/earned-exposure/edge-severance-events-three-valued-survival)
- [Emitting the Deference Correction Linkage Record to the Followed Agent: Reporting a Downstream Correction Without Producing a Rating \(/articles/earned-exposure/emission-deference-correction-linkage-record-followed\)](/articles/earned-exposure/emission-deference-correction-linkage-record-followed)
- [The No-Address Omission Branch: What a Deferring Agent Records When an Exposure Dispatch Reaches Nobody \(/articles/earned-exposure/exposure-dispatch-no-address-omission-branch\)](/articles/earned-exposure/exposure-dispatch-no-address-omission-branch)
- [Channel-Baseline Normalization: Taking Channel Overhead Out of the Latency Differential \(/articles/earned-exposure/latency-differential-channel-baseline-normalization\)](/articles/earned-exposure/latency-differential-channel-baseline-normalization)

- [Reason-Type Resolution From Severance-Survival States: Typing an Agent's Relationship Edges Without Asking the Counterparty](/articles/earned-exposure/reason-type-resolution-survival-states) (/articles/earned-exposure/reason-type-resolution-survival-states).
- [Reason-Type Retyping by a Superseding-Type Entry: Recomputing an Edge's Reason-Type on Later Severance Evidence Without Disturbing Modifications Already Applied](/articles/earned-exposure/reason-type-retyping-superseding-type-entry) (/articles/earned-exposure/reason-type-retyping-superseding-type-entry).
- [Renewal-Side Replenishment of the Authorization Budget: Restoring an Agent's Metered Spending Capacity Only on Exposure to an Uncounted Origin-Equivalence Class](/articles/earned-exposure/renewal-side-replenishment-authorization-budget) (/articles/earned-exposure/renewal-side-replenishment-authorization-budget).
- [Following Does Not Count as Being Tested: Retention of the Short Decay Constant When an Agent Defers, and the Action Class That Stays Untested](/articles/earned-exposure/retention-short-decay-constant-deference-class) (/articles/earned-exposure/retention-short-decay-constant-deference-class).
- [Standing-Determination Designation: Substituting a Verified Counterparty Determination for One Conjunct of the Dispatch-Authority Predicate](/articles/earned-exposure/standing-determination-designation-substitution-inside-dispatch) (/articles/earned-exposure/standing-determination-designation-substitution-inside-dispatch).
- [Uncounted-Class Renewal: Gating Agent Authorization on a Windowed Register of Origin-Equivalence Classes](/articles/earned-exposure/uncounted-class-renewal-gated-windowed-register) (/articles/earned-exposure/uncounted-class-renewal-gated-windowed-register).

APPLICATIONS • GENERAL

- [Vendor Qualification Without Ratings: Standing Earned From Recorded Conduct](/articles/earned-exposure/vendor-qualification-without-ratings) (/articles/earned-exposure/vendor-qualification-without-ratings).
- [Credentialing Autonomous Practitioners: Deference That Survives Across Agents](/articles/earned-exposure/professional-credentialing-agents) (/articles/earned-exposure/professional-credentialing-agents).

APPLICATIONS • SPECIFIC

- [Marketplace Seller Reputation Systems: Presented Scores Versus Earned Admissibility](/articles/earned-exposure/marketplace-seller-reputation) (/articles/earned-exposure/marketplace-seller-reputation).
- [Letta: Agent Memory and Standing Earned From Conduct](/articles/earned-exposure/letta-agent-memory) (/articles/earned-exposure/letta-agent-memory).
- [Trustpilot Signals and Conduct-Derived Admissibility](/articles/earned-exposure/trustpilot-signals) (/articles/earned-exposure/trustpilot-signals).
- [Copilot Coding Agents: Review Gates and Earned Exposure](/articles/earned-exposure/github-copilot-agent) (/articles/earned-exposure/github-copilot-agent).
- [Supplier Risk Scores and Recorded Counterparty Conduct](/articles/earned-exposure/supplier-risk-scoring) (/articles/earned-exposure/supplier-risk-scoring).
- [Verified Profiles and Standing Built From Recorded Conduct](/articles/earned-exposure/linkedin-verification) (/articles/earned-exposure/linkedin-verification).

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 → \(/earned-exposure\)](#)