Nick Clark · Published July 23, 2026 · PDF

The Valuation-Free Ordinal Standing Quantity

When a governed AI agent must decide how heavily an incoming correction should weigh on its own conduct, the obvious lever is the counterparty's standing, and the obvious way to measure standing is a cardinal score. The valuation-free ordinal standing quantity, disclosed in U.S. Provisional Application No. 64/117,812, refuses that lever. It derives a bounded standing multiplier from three timing-domain statistics computed over the agent's own recorded history with one counterparty: the direction of initiation, the latency to reciprocate, and the regularity of exchange. Each statistic maps through two declared breakpoints to a score of zero, one, or two, the three scores sum to an integer from zero to six, and a declared non-decreasing map turns that integer into the multiplier. No cardinal magnitude of benefit is ever computed, solicited, or stored.

1. The Gap

An autonomous agent that accepts correction has to decide how much any single correction is worth. Not every incoming evaluation should move the agent's internal state by the same amount. A correction from a party the agent has transacted with steadily for a year is not obviously the same as one from a party it has never dealt with, and most systems that confront this reach for a familiar tool: a standing score attached to the counterparty, used to scale the weight.

The trouble is where that score comes from. A cardinal standing score is a magnitude, and a magnitude has to be sourced. Either the counterparty supplies it, in which case the party being weighed controls its own weight, or a central rater assigns it, in which case the agent has handed that judgment to a third party, or it is derived from the value of past dealings, in which case a well-funded counterparty can buy its way up. Each path shares one defect: once the system computes or stores a magnitude of benefit conferred or received, standing behaves like a price, and a price can be bid up.

The valuation-free ordinal standing quantity closes this gap by declining to compute a magnitude at all. It produces a bounded multiplier that scales how much a correction moves the agent's own conduct state, and it derives that multiplier entirely from the shape of the agent's own recorded exchange with the counterparty in the time domain. Nothing a counterparty declares enters it. No worth, price, rating, or valuation is read, requested, or retained.

2. Mechanism

The standing quantity operates inside reason-type-weighted admission. When an accepted determination arises from an asserting party, a state modifier changes the agent's scoped integrity vector and its self-esteem aggregate by some magnitude. That magnitude is scaled by an **effective weighting**, and the effective weighting is the product of two factors: a per-reason-type weighting derived from the severance-survival character of the relationship, and a **standing multiplier**. The standing quantity is the machinery behind that second factor.

The multiplier is derived from three timing-domain statistics computed over the agent's own append-only lineage entries for the counterparty, and from nothing else.

Direction of initiation. The first statistic is equal to the count of interactions the agent itself initiated divided by the total count of interactions with that counterparty. It captures which side reaches out. A relationship in which one party does all the

initiating has a different shape from one in which initiation runs both ways, and the direction statistic reads that shape from the record of who opened each exchange.

Latency to reciprocate. The second statistic is the median reciprocation interval between the parties: the typical gap between an interaction initiated by one party and the next interaction initiated by the other. It measures how promptly exchange comes back around. The interval is not measured in wall-clock time. Each interval is counted in successor epochs of the agent's own dynamic hash chain, so latency is expressed in the agent's own advancing state rather than against a wall-clock constant.

Regularity. The third statistic is the median absolute deviation of the successive inter-interaction intervals. It measures how evenly spaced the exchange is. A low value indicates steady, regular exchange, and a high value indicates exchange that arrives in bursts separated by long gaps. Like latency, regularity is computed over intervals counted in successor epochs of the hash chain.

Each of the three statistics is then reduced to a small ordinal score. Each maps through **two declared breakpoints** to a component score of zero, one, or two. The breakpoints are carried in the signed policy object, so the policy, not the counterparty and not any outside rater, fixes where a statistic falls from a higher score to a lower one. Three statistics, each yielding zero, one, or two, sum to an integer **standing quantity in the range zero to six inclusive**.

That integer becomes the multiplier through a declared non-decreasing map. Two properties of the map matter. First, it is non-decreasing in the standing quantity, so a higher standing quantity never yields a smaller multiplier than a lower one. Second, it is bounded at both ends: the map is declared with a least multiplier strictly above zero and a declared greatest multiplier. The lower bound means a low standing quantity cannot drive the multiplier to zero, so a counterparty with the weakest recorded history

still moves the agent's state, only by less. The upper bound means a strong history cannot amplify a correction without limit. Standing can tilt the weight; it can neither silence a party nor let one dominate.

The defining negative limitation runs through every step. No cardinal magnitude of benefit conferred or received between the agent and the counterparty is computed, solicited, or stored. There is no price on an interaction, no valuation of the relationship, no rating drawn from the counterparty, and no scalar representing worth. Every input is a count, an ordering, a direction, or a timing already present in entries the agent holds in its own lineage, combined only with breakpoints the agent's policy declares. The output is ordinal by construction, and it is derived from the timing and direction of recorded exchange alone.

3. Operating Parameters

The filed disclosure declares the following, and the article states no numeric values beyond what it declares.

- **Component score domain.** Each statistic yields a component score of zero, one, or two.
- **Breakpoints.** Each of the three statistics maps to its score through two breakpoints declared in the signed policy object, for six declared breakpoints in all.
- **Standing quantity domain.** The three component scores sum to an integer standing quantity from zero to six inclusive, spanning seven admissible values.
- **The multiplier map.** A non-decreasing map, declared in the signed policy object, carries the standing quantity to the standing multiplier. It is declared with a least multiplier strictly greater than zero and a declared greatest multiplier, so the multiplier is bounded away from zero and above.
- **Interval units.** Every interval, for both latency and regularity, is counted in successor epochs of the agent's dynamic hash chain rather than in wall-clock time.

- **Effective weighting.** The standing multiplier enters admission as one factor of a product; the other factor is the per-reason-type weighting. Where the reason-type is not-typeable and that product falls below a declared non-zero minimum, the effective weighting is raised to that minimum, the floor applying to the product rather than to the weighting alone.

Relatedly, the observation windows the surrounding architecture uses are keyed to the agent's own recorded arrival rate through a recorded inter-arrival estimator, a declared measure of central tendency over intervals between successive lineage entries, so those windows bind to the agent's own history rather than to a fixed external constant.

4. Composition

The standing quantity is one half of the weight a correction carries, and it is designed to sit beside the other half rather than replace it. Reason-type-weighted admission first resolves the reason-type of the edge to the asserting party from recorded severance-survival states, then retrieves a per-reason-type weighting. The standing multiplier multiplies that weighting to form the effective weighting applied to the scoped integrity vector and the self-esteem aggregate. Two independent readings of the relationship, one drawn from its severance history and one from its timing, therefore combine into a single magnitude.

The not-typeable floor knits the two together. When the edge has not been severance-tested, the per-reason-type weighting is at its lowest, and a low standing multiplier could in principle drive the combined product toward nothing. The floor on the product prevents that, and it works in concert with the standing map's least-multiplier-above-zero property: neither factor, alone or together, can reduce the movement of state to zero. A newly encountered or low-standing counterparty still corrects the agent, which is the property that keeps the agent correctable by parties it has no history with.

Both factors share a discipline that runs across the whole filing: they are computed from the agent's own records, never from anything the asserting party declares. The reason-type is not declared by the asserting party, and the standing quantity reads only the agent's own lineage. This is the same principle the architecture applies wherever agents meet, where no rating, score, standing quantity, or credential passes between two agents. In the filing's defined terms, a standing quantity is an ordinal quantity computed from a direction of initiation, a latency to reciprocate, and a regularity of recorded interactions, with no cardinal magnitude of benefit computed or stored. And it is the only quantity the architecture will treat as one: elsewhere the filing bars any count of corrections or adoptions an agent holds for another party from being consumed as a standing quantity in an admission or dispatch decision. This ordinal figure is therefore not one reputation signal among several but the single form standing is permitted to take. Its inputs, being counts and successor-epoch intervals in the append-only lineage, are the same primitives the counterparty identity record and the dynamic hash chain supply throughout the architecture, so it composes without introducing any new external quantity.

5. Prior-Art Distinction

The problem of weighting one party's input by that party's standing is old, and several well-developed families of technique address it. Each is distinguished here structurally, not by any suggestion that a given system trespasses on this disclosure.

Reputation and trust-score systems assign a party a cardinal scalar, typically aggregated from third-party ratings or from the recorded outcomes of past transactions. The structural distinction is at the input: such a scalar is a magnitude imported from raters or from valued outcomes, whereas the standing quantity computes no magnitude of benefit and reads no rating supplied by any party, taking only the timing and direction of the agent's own recorded exchange.

Recommender and collaborative-filtering weightings weight a source by a predicted value or affinity, which is again a cardinal estimate the system computes and stores. The standing quantity predicts no value; it emits an integer from zero to six and a bounded multiplier, and it stores no estimate of worth.

Graph-centrality and link-based trust, in the style of network-authority ranking, derive standing from a party's position in a global graph of endorsements or links. That approach requires a network and third-party edges. The standing quantity is dyadic and local: it is computed from one edge's timing in the agent's own lineage, needs no global graph, and consults no third party's endorsement.

Recency-frequency-monetary scoring, familiar from commercial segmentation, is the closest surface analogue, since its recency and frequency components echo latency and regularity. The decisive difference is the monetary component, which is exactly the cardinal magnitude of benefit this mechanism excludes, and the purpose: such scoring produces a value estimate for targeting, while the standing quantity produces only a bounded, non-decreasing multiplier that scales how far a correction moves the agent's own integrity state and can never extinguish it.

Across all four families, the common feature this mechanism removes is the cardinal magnitude. Where the prior categories compute or store a number standing for value, worth, or predicted benefit, the standing quantity is built to declare, at every step, that no such number is computed, solicited, or stored.

6. Disclosure Scope

The operative disclosure is U.S. Provisional Application No. 64/117,812. The valuation-free ordinal standing quantity is disclosed at Section 10.1, paragraph [0336], with the reason-type weighting into which it feeds at paragraph [0335], the recorded inter-arrival estimator at paragraph [0337], and the definition of "standing quantity" at paragraph [0541]. What is disclosed there and asserted here is: derivation of a standing

multiplier from three timing-domain statistics over the agent's own lineage, namely direction of initiation equal to agent-initiated over total interactions, a latency equal to the median reciprocation interval, and a regularity equal to the median absolute deviation of successive inter-interaction intervals, each interval counted in successor epochs of the agent's hash chain; the mapping of each statistic through two declared breakpoints to a component score of zero, one, or two; the summation to an integer standing quantity from zero to six; the declared non-decreasing map to a multiplier with a least value above zero and a declared greatest value; and the express non-computation of any cardinal magnitude of benefit. This article states no particular numeric values for the breakpoints, the multiplier bounds, or the map, and it does not reproduce every sub-embodiment appearing in the inventor's fuller exposition. Nothing here should be read to enlarge or narrow the provisional's disclosure.

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)