

# A Typed, Dimensionless Conduct Divergence Measure

A persistent semantic agent that acts under different declared policies in different scope partitions will naturally behave differently across them, and any consistency check that reads that difference as misconduct punishes the agent for the differentiation its operator declared. The typed, dimensionless conduct divergence measure, disclosed in U.S. Provisional Application No. 64/117,812, is the measurement stage that makes such a check tractable without a language model. The agent retrieves executed-action entries from two scope partitions over a declared comparison window, forms comparison dimensions as action-class and affected-party-class pairs that each partition populates at a declared minimum count, and computes one divergence component per dimension from recorded outcome fields alone. A type designator on each outcome field selects a proportion, rank-distribution, or scaled central-tendency difference, so every component is a pure number and the components are mutually comparable.

---

## 1. The Gap

An autonomous agent that acts under one declared policy in one scope partition and a different declared policy in another will act differently across the two. That is not a defect. It is the point of letting an operator declare separate values for separate scopes.

The difficulty is that a consistency check run over the raw difference between the two bodies of conduct cannot tell two things apart: the differentiation an operator deliberately declared for one scope, and conduct that diverges across scopes in a respect no declared policy distinguishes. Read every cross-scope difference as misconduct and you punish declared differentiation; wave all of it away as context and genuine cross-scope inconsistency goes undetected.

One conventional way to ask whether two bodies of conduct diverge is to give the material to a language model and take back a judgment. A single such judgment is not computed from the recorded outcome fields, and it collapses outcomes of different kinds into one assessment. A later stage is then left with nothing structured to work with: no per-dimension, per-outcome-type quantity from which the divergence a declared policy explains can be separated from the divergence it does not.

The mechanism disclosed here is the measurement primitive. It produces a divergence quantity between two scope partitions, computed from recorded outcome fields alone, typed to the kind of outcome each field records, and rendered dimensionless so that components drawn from enumerated, ordinal, and interval-valued outcomes are mutually comparable. It performs no language-model inference, and it appends its result together with the comparison window and the entries relied upon to the append-only lineage field, so the same components can be computed again from the record.

## **2. Mechanism**

The agent computes a conduct divergence measure between two scope partitions. The computation is an ordered procedure over recorded conduct.

**Retrieval over a declared window.** The agent retrieves, from each of the two compared scope partitions, entries recording executed actions that fall within a comparison window declared in the signed policy object. Each retrieved entry carries an

action-class identifier, an affected-party class, a recorded outcome field, and a recorded time. Only executed-action entries are drawn, so the measure runs over conduct the agent carried out rather than over any account of it.

**Formation of comparison dimensions.** The agent forms comparison dimensions, each a pair of an action-class identifier and an affected-party class. A pair becomes a comparison dimension only where both compared partitions hold retrieved entries for that pair at the declared minimum count. A pair that either partition populates below that count forms no comparison dimension and contributes nothing.

**Typed per-dimension components.** For each comparison dimension, the agent computes a single per-dimension divergence component from the recorded outcome fields of the two partitions on that dimension, and from those fields alone. The form of the component is chosen by a type designator carried on the outcome field, which selects exactly one of three types.

- Where the outcome field is of the enumerated type, the component is a proportion difference: the difference between the two partitions' proportions of entries recording each enumerated outcome value.
- Where the outcome field is of the ordinal type, the component is a rank-distribution difference between the two partitions over the ordinal levels of that field.
- Where the outcome field is of the interval-valued type, the component is a central-tendency difference: the difference between a central-tendency statistic of the first partition's outcomes and the same statistic of the second partition's, divided by a per-dimension divergence scale declared for that dimension.

Because the type designator selects the component form per outcome field, one procedure yields a component appropriate to each field, whether the outcomes it records are enumerated, ordinal, or interval-valued.

**Dimensionlessness and mutual comparability.** A proportion difference and a rank-distribution difference are dimensionless as computed. An interval-valued difference is rendered dimensionless by the division by the per-dimension divergence scale, which carries the units of that outcome field out of the quotient. Every component therefore sits on a common scale, whatever outcome type it came from. That common scale is what lets a downstream stage compare components drawn from enumerated, ordinal, and interval-valued outcomes and test the residual against a declared residual bound.

**No language-model inference.** No step reads natural-language content, and no step invokes a language model to classify or score the conduct. The measure is arithmetic over recorded fields alone.

**The measure and its record.** The conduct divergence measure is the vector of per-dimension divergence components over the comparison dimensions that were formed, appended to the append-only lineage field together with the comparison window and the entries relied upon. Because the vector and its inputs are on the record, a reviewer can retrieve the same window and the same entries and compute the same components. The measure is a vector rather than a scalar so that a later stage can act on individual dimensions, holding the dimensions where two scopes diverge apart from the dimensions where they agree.

The same procedure is disclosed for the case of two persistent semantic agents in place of two partitions of one agent, each agent resolving a policy object under a common alias and the presenting agent's entries admitted only on signature and hash-chain succession, so that the identical typed, dimensionless construction measures cross-agent consistency as well as an agent's consistency with itself.

### 3. Operating Parameters

Every parameter of the measure is carried in the signed policy object, and the filed disclosure states them as declared quantities rather than fixed constants.

- **Comparison window.** The window over which executed-action entries are retrieved is declared in the signed policy object.
- **Minimum count.** The minimum number of retrieved entries a partition must hold on an action-class and affected-party-class pair for that pair to form a comparison dimension is declared in the signed policy object, and the gate is applied to each compared partition independently.
- **Per-dimension divergence scale.** For an interval-valued outcome field, the scale by which the central-tendency difference is divided is declared per comparison dimension in the signed policy object. It is what carries the units out of the component and leaves a dimensionless number.
- **Outcome-field type designator.** Each outcome field carries a type designator, declared for the action class, that selects exactly one member of a closed set of three: enumerated, ordinal, or interval-valued. The designator determines which component form is computed.

The disclosure fixes no numeric default for any of these. The window, the minimum count, the scale, and the type are the operator's to declare, and the measure computes over whatever the policy object states.

### 4. Composition

The measure is the first stage of the residualized-divergence procedure, and at this stage it consults no policy. It reports how far two scopes' recorded outcomes differ, without yet distinguishing legitimate difference from the rest. The stages that follow do the separating. The agent identifies the policy object that governed each partition at the time of the compared conduct, computes a policy-difference set describing where the

two governing policies differ, and removes the component of divergence those declared differences explain. What remains is the residual: the divergence in respects the declared policies do not distinguish.

That division of labor is why the measure is built the way it is. Because each component is typed and dimensionless, the policy-difference stage can subtract a difference of declared bounds from a component on the same common scale, dimension by dimension, and retain only any excess. And because the measure is a vector rather than a scalar, the residual can be confined to the specific dimensions on which the divergence is unexplained rather than smeared across the whole comparison.

Only the residual is consumed by consistency evaluation. The component a policy difference explains moves no value of the scoped integrity vector, moves no value of the self-esteem aggregate, and increments no counter. Where the residual exceeds a declared bound on at least one dimension, the agent modifies the scoped integrity vector and the self-esteem aggregate of each compared partition where the diverging conduct is recorded, recomputes each partition's deviation quantity, and writes the authorization gate to the withheld state confined to the partition concerned. The gate write stays within the partition where the diverging conduct is recorded rather than propagating across partitions, so the measure composes with the integrity and deviation machinery that governs permitted deviation without spreading a single scope's finding to the others.

Feeding the residual, this measure is what surfaces divergence between scopes that no declared policy difference accounts for, and it reaches that finding without any step soliciting a justification or consulting a party.

## **5. Prior-Art Distinction**

Several established categories measure something adjacent, and the measure is structurally distinct from each.

**Language-model evaluation of agent behavior.** A common contemporary approach scores an agent's conduct, or judges the consistency of two bodies of conduct, by prompting a model to read the material and return an assessment. The measure disclosed here performs no language-model inference of any kind. It computes over recorded outcome fields and appends its result with the window and the entries relied upon, so the same components can be computed again from the record; a model that reads the material and returns a verdict does neither.

**Scalar distribution-distance and drift metrics.** Monitoring and machine-learning systems compare two distributions with a single distance, such as a divergence between probability distributions or a population-stability index, typically assuming one kind of variate and returning one number. The measure here instead selects a type-appropriate component per outcome field, enumerated, ordinal, or interval-valued, and returns a vector of dimensionless components indexed by action-class and affected-party-class dimension. It is the typing and the per-dimension structure, not a single aggregate distance, that carry the work.

**Fairness and disparate-impact auditing.** Bias audits compare outcomes across populations, often with a fixed metric per protected attribute. The comparison here is between an agent's own two scope partitions, or between two agents under a common alias, and is a self-consistency measurement rather than a cross-population one. Its components are selected by outcome type and placed on a common scale so a later stage can subtract declared policy differences, a step that has no counterpart in group-fairness metrics.

**Policy-compliance and rule-checking engines.** Compliance tooling tests conduct against a rulebook and reports violations. At its measurement stage this measure consults no rulebook: it reports typed divergence, and the separation of legitimate from illegitimate difference is performed by distinct downstream stages that subtract what the governing policies explain. The measurement primitive and the policy comparison are kept apart.

None of these categories is disparaged here, and nothing in this article asserts or implies that any product or party practicing them infringes anything. The distinction is structural.

## 6. Disclosure Scope

The operative, citable disclosure is U.S. Provisional Application No. 64/117,812, Section 10.3, paragraph [0375], which discloses the typed, dimensionless conduct divergence measure: retrieval of executed-action entries over a declared comparison window; formation of comparison dimensions as action-class and affected-party-class pairs each partition holds at a declared minimum count; computation of a per-dimension divergence component from recorded outcome fields alone; selection among a proportion difference, a rank-distribution difference, and a central-tendency difference divided by a declared per-dimension divergence scale according to an outcome-field type designator naming enumerated, ordinal, or interval-valued; the rendering of every component dimensionless and mutually comparable; the absence of any language-model inference; and the appending of the resulting vector to the append-only lineage field with the window and the entries relied upon.

This article asserts no numeric value for any parameter beyond stating that the comparison window, the minimum count, the per-dimension divergence scale, and the outcome-field type designator are each declared in the signed policy object. Specific numeric floors, window bounds, worked numeric examples, and any illustrative parameter magnitudes are not disclosed here and are reserved to the applicant. Nothing in this article enlarges, narrows, or supersedes the filed provisional, which controls.

Bounded, accountable deviation from declared values — measured, not forbidden.

Provisional application

## PRIMARY TECHNICAL DISCLOSURE

- [Permitted Deviation: Bounded, Recorded Departure From an Agent's Declared Values \(/articles/permitted-deviation/bounded-departure-from-declared-values\)](/articles/permitted-deviation/bounded-departure-from-declared-values).

## SECONDARY TECHNICAL

- [A Typed, Dimensionless Conduct Divergence Measure \(/articles/permitted-deviation/typed-dimensionless-divergence-measure\)](/articles/permitted-deviation/typed-dimensionless-divergence-measure).
- [The Policy-Difference Set and Policy-Explained Divergence \(/articles/permitted-deviation/policy-explained-divergence\)](/articles/permitted-deviation/policy-explained-divergence).
- [Residual Divergence and Its Confined Consumption \(/articles/permitted-deviation/residual-divergence-confined-consumption\)](/articles/permitted-deviation/residual-divergence-confined-consumption).
- [Propagation Halt on Undischarged Reparation Arcs \(/articles/permitted-deviation/propagation-halt-undischarged-arcs\)](/articles/permitted-deviation/propagation-halt-undischarged-arcs).
- [Attack-Conditional Propagation Suspension \(/articles/permitted-deviation/attack-conditional-propagation-suspension\)](/articles/permitted-deviation/attack-conditional-propagation-suspension).

## TERMINOLOGY

- [aggregate retention \(/glossary#aggregate-retention\)](/glossary#aggregate-retention).
- [deviation deductible \(/glossary#deviation-deductible\)](/glossary#deviation-deductible).
- [deviation likelihood \(/glossary#deviation-likelihood\)](/glossary#deviation-likelihood).
- [entropy-weighted harm coefficient \(/glossary#entropy-weighted-harm-coefficient\)](/glossary#entropy-weighted-harm-coefficient).
- [need quantity \(/glossary#need-quantity\)](/glossary#need-quantity).
- [permitted deviation \(/glossary#permitted-deviation\)](/glossary#permitted-deviation).
- [self-esteem aggregate \(/glossary#self-esteem-aggregate\)](/glossary#self-esteem-aggregate).

---

[Integrity Quantities and Permitted Deviation overview → \(/permitted-deviation\)](/permitted-deviation).