



[Home](#) [Licensing](#) [Patents](#) [Articles](#)

Integrity-Aware Multi-Agent Negotiation

by [Nick Clark](#) | Published March 27, 2026 | [PDF](#)

Integrity state of each agent influencing weight given to contributions, votes, and proposals in multi-agent collaboration and collective decision-making.

What It Is

Integrity state of each agent influencing weight given to contributions, votes, and proposals in multi-agent collaboration and collective decision-making.. This mechanism is defined in Chapter 3 of the cognition patent as a structural component of the agent's cognitive architecture, operating through deterministic evaluation rather than heuristic approximation.

Every aspect of this mechanism is specified declaratively in the agent's policy reference, making it auditable, reproducible, and governable without requiring access to the agent's internal decision-making process.

Why It Matters

Without integrity-aware multi-agent negotiation, agents lack the structural mechanism needed to maintain behavioral coherence across changing conditions. Existing systems either enforce rigid rules that cannot adapt to context, or rely on probabilistic heuristics that provide no guarantees about normative consistency. The gap is between systems that are rule-following and systems that are genuinely coherent.

The consequence of this gap appears in multi-agent coordination, governance compliance, and long-running autonomous operation. Agents that cannot track their own normative consistency accumulate behavioral drift over time, gradually deviating from their declared values without any mechanism to detect or correct the deviation.

How It Works Structurally

As defined in Chapter 3 of the cognition patent, integrity-aware multi-agent negotiation operates through a deterministic evaluation function embedded within the agent's cognitive architecture. The function receives structured inputs from the agent's canonical fields and produces outputs that govern subsequent processing stages. Every input, computation step, and output is recorded in the agent's lineage, ensuring complete reproducibility.

The implementation maintains its own state within the agent's integrity field, which persists across execution cycles and substrate migrations. Policy constraints govern every parameter, threshold, and behavioral boundary. Cross-primitive coupling ensures that changes in the integrity field propagate to confidence governance, forecasting, and discovery traversal through defined interfaces.

What It Enables

This mechanism enables agents that maintain verifiable behavioral coherence over extended operational periods. Governance bodies can audit not just individual decisions but the complete integrity trajectory that led to them. Multi-agent systems gain a structured basis for evaluating the trustworthiness of collaborating agents based on their demonstrated integrity.

Because this mechanism is policy-governed and deterministic, it can be formally analyzed, audited, and certified. Regulatory compliance is demonstrable through structural analysis rather than solely through empirical testing. Different domains can tune the mechanism's parameters through policy configuration without requiring architectural changes, making the same structural capability applicable to autonomous vehicles, companion AI, therapeutic agents, and enterprise systems.

[Integrity & Coherence All 21 steps →](#)

Track normative consistency. Detect deviation. Self-correct.

Primary Technical Disclosure

◦ [The Coherence Trifecta: Empathy, Integrity, and Self-Esteem as a Unified Control Loop](#)

Secondary Technical

◦ [Coping Under Empathic Pressure: HSP, Narcissism, and Psychopathy as Control-Loop Intercepts](#) ◦ [Three-Domain Integrity Model](#) ◦ [Deviation Function \$D=\(N-T\)/\(ExS\)\$](#) ◦ [Self-Esteem as Internal Validator](#) ◦ [Deviation as Deterministic Semantic Mutation](#) ◦ [Integrity Structural Placement](#) ◦ [Empathy as Distributed Moral Load](#) ◦ [Coherence Trifecta Control Loop](#) ◦ [Coping Intercept Patterns](#) ◦ [Integrity Deviation Logging](#) ◦ [Integrity Collapse Detection](#) ◦ [Redemption Engine](#) ◦ [Moral Trajectory Forecasting](#) ◦ [Integrity-Aware Trust Slope Validation](#) ◦ [Integrity-Confidence Cross-Primitive Coupling](#) ◦ [Integrity-Modulated Discovery Traversal](#) • [Integrity-Aware Multi-Agent Negotiation](#) ◦ [Biological Signal Coupling for Integrity](#) ◦ [Policy-Based Integrity Constraints](#) ◦ [Integrity Field Portability](#) ◦ [Predictive Deviation Alerting](#) ◦ [Governed Forgetting](#) ◦ [Predictive Social Modeling](#)

Applications (General)

◦ [Autonomous Vehicle Ethical Decision-Making Through Computable Integrity](#) ◦ [Financial Trading Systems That Track Their Own Normative Consistency](#) ◦ [Integrity and Coherence for Legal Advisory Agents](#) ◦ [Integrity and Coherence for Government Policy Agents](#) ◦ [Integrity and Coherence for Journalism Editorial Agents](#) ◦ [Integrity and Coherence for Environmental Compliance Agents](#) ◦ [Integrity and Coherence for Insurance Underwriting Agents](#) ◦ [Integrity and Coherence for Social Media Moderation Agents](#)

Applications (Specific)

◦ [Waymo's Ethical Decisions Have No Normative Memory](#) ◦ [Cruise's Safety System Cannot Track Its Own Consistency](#) ◦ [JPMorgan's Trading Compliance Has No Normative Trajectory](#) ◦ [Palantir's Analytics Cannot Monitor Their Own Normative Drift](#) ◦ [Aurora's Self-Driving Stack Has No Normative Memory](#) ◦ [Nuro's Delivery Robots Optimize Without Normative Tracking](#) ◦ [Zoox Plans Maneuvers Without Tracking Normative Drift](#) ◦ [Motional Validates Safety Without Governing Normative Trajectory](#) ◦ [Argo AI's Shutdown Reveals the Cost of Missing Normative Architecture](#) ◦ [comma.ai Learns to Drive Without Learning Ethics](#)

[Integrity & Coherence overview →](#)

AQ

deterministic
autonomy

Legal

Subject to one or more pending U.S. and international patent applications, see [Patents](#) for the current list and status. No license, express or implied, is granted. Any use requires a separate written agreement—see [Licensing](#). Patent applications referenced on this site are pending. Claim scope, if any, is subject to examination and may issue in altered form or not at all. See [Legal](#) for terms and conditions.

Adaptive Query™ is a trademark of Nicholas Clark. U.S. federal registration is pending. federal registration. AQ™, AQ Inside™, Adaptive Index™, Adaptive Network™, Semantic Agent™, @AQ™, AQID™, and Adaptive Coin™ are used as trademarks in connection with the Adaptive Query platform and brand. Other names may be trademarks of their respective owners.

Platform operated by Adaptive Query LLC, which provides patent and trademark licensing services. Copyright © 2025-2026 Nicholas Clark. All rights reserved.

Last updated: 2026-03-03



-
- [Inventive Steps](#)
- [Licensing](#)
- [Patents](#)
- [Articles](#)
- [Legal](#)
- [Opportunities](#)
- [Sitemap](#)



-
- nick@qu3ry.net
- 72 28 14 36 01



[Invented by Nick Clark](#) | Founding Investors: Devin Wilkie