



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

Waymo's Ethical Decisions Have No Normative Memory

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

Waymo operates the most mature autonomous driving fleet in public service. Its perception, prediction, and planning stack handles millions of miles of real-world driving with a safety record that exceeds human performance. But when the system faces ethical edge cases, each scenario is evaluated independently against predefined rules. The vehicle has no persistent normative state, no memory of its own ethical trajectory, and no mechanism to detect drift in its decision-making consistency over time. Resolving this requires integrity as a first-class cognitive primitive.

What Waymo built

Waymo's engineering achievement is substantial. The system integrates lidar, camera, and radar perception with behavioral prediction models that anticipate the actions of other road users, and a planning stack that generates safe trajectories in real time. The safety framework includes multiple layers

of redundancy, conservative behavior policies, and extensive simulation testing. Millions of autonomous miles in public service demonstrate that the system works.

Ethical edge cases, situations where the vehicle must choose between competing safety objectives, are handled through a hierarchy of rules. Protect occupants. Protect pedestrians. Minimize harm. Obey traffic law. These rules are encoded in the planning system's cost functions and constraint sets. When a conflict arises, the system evaluates the specific scenario against these rules and selects the action that best satisfies the hierarchy.

The gap between per-scenario rules and normative consistency

The structural limitation is not in any single ethical decision. It is in the relationship between decisions over time. A vehicle that operates for years makes thousands of decisions with ethical dimensions. Each intersection where a pedestrian is near the crosswalk. Each merge where yielding creates a different risk profile than proceeding. Each school zone where the balance between traffic flow and safety margin must be recalibrated.

These decisions, taken individually, may each satisfy the rule hierarchy. Taken collectively, they constitute an ethical trajectory. A vehicle that has gradually become more aggressive in yielding decisions, not because any single decision violated a rule but because the cumulative pattern shifted, has experienced normative drift. Current architecture has no mechanism to detect this drift because there is no persistent normative state to drift from.

The deviation function $D=(N-T)/(E \times S)$, where N is current normative position, T is the agent's established trajectory, E is empathy weighting, and S is self-esteem stability, provides a computable measure of how far the agent's current behavior has diverged from its own normative baseline. Without this function, the vehicle cannot distinguish between a single unusual decision and a systematic shift in its ethical posture.

Why rule hierarchies are not enough

Rule hierarchies define what is permissible. They do not track what has been chosen. A vehicle that consistently selects the minimum-safety-margin option in ambiguous scenarios is operating within its rules but is making a different ethical statement than one that consistently selects the maximum-safety-margin option. Both are rule-compliant. They represent different normative positions. Without persistent integrity state, neither vehicle knows which position it has been occupying.

The problem intensifies with fleet-wide consistency. If one Waymo vehicle in San Francisco has drifted toward aggressive merging while another has drifted toward conservative behavior, the fleet presents inconsistent ethical behavior to the public. Passengers in different vehicles experience different normative standards. The fleet has no structural mechanism to detect or correct this divergence because no individual vehicle maintains normative state that can be compared across the fleet.

What integrity coherence enables

With integrity as a first-class cognitive primitive, each vehicle maintains a three-domain integrity model: behavioral integrity tracking consistency of actions, normative integrity tracking alignment with declared values, and narrative integrity tracking coherence of the vehicle's operational story over time. The deviation function computes the distance between current behavior and established baseline continuously.

When deviation exceeds a threshold, the coherence trifecta engages. The vehicle's empathy mechanism considers how its decisions affect other road users. Its self-esteem validator checks whether the current behavioral pattern is consistent with its defined role. Its integrity tracker determines whether a coping intercept is needed: a structural correction that returns the vehicle to its normative baseline without requiring external intervention.

This gives Waymo something its current architecture cannot provide: a vehicle that not only follows rules but monitors the consistency of its own rule-following, detects when its behavioral pattern is drifting, and self-corrects before the drift becomes operationally significant.

The structural requirement

Waymo's per-scenario ethical evaluation is sound. The gap is longitudinal. An autonomous vehicle that operates for years in complex environments needs normative memory: persistent state that tracks what it has chosen, how those choices form a pattern, and whether that pattern remains consistent with its declared ethical framework. This is not achievable through more rules or better simulation. It requires integrity as a persistent, computed cognitive state with deviation tracking and self-correction.

[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=\(N-T\)/\(E \times S\)\$](#) [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