



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

## Confidence Governance for Pharmaceutical Dosing Systems

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

Medication dosing errors are among the most common causes of preventable patient harm. AI dosing systems that recommend drug doses based on patient data must handle conflicting lab values, incomplete records, drug interactions, and patient-specific factors. Current systems generate recommendations with stated confidence intervals but continue recommending regardless of how uncertain the inputs are. Confidence governance provides risk-proportional thresholds that require higher confidence for higher-risk medications and a non-executing mode that pauses dosing recommendations when clinical confidence falls below the safety threshold for the specific drug and patient context.

---

**The confidence gap in clinical dosing**

Clinical dosing decisions depend on multiple data inputs: patient weight, organ function, lab values, concurrent medications, genetic factors, and clinical context. Each input carries its own uncertainty. Lab values may be from samples taken hours ago. Organ function assessments may be based on estimates rather than direct measurement. Drug interaction databases may not cover the patient's specific medication combination.

Current AI dosing systems process these inputs and generate a recommended dose. The recommendation may include a confidence interval, but the system does not modulate its behavior based on confidence. A dosing recommendation generated with high confidence from complete, recent data looks the same to the clinician as one generated with low confidence from incomplete, conflicting data. Both appear as a specific dose recommendation on the screen.

The consequence is that clinicians must independently assess the reliability of each recommendation, a cognitive task that is particularly error-prone during high-workload situations when AI dosing assistance is most needed. The system generates the recommendation. The clinician must determine whether to trust it. The system provides no structural help with the trust assessment.

## Risk-proportional confidence thresholds

Confidence governance introduces risk-proportional thresholds: medications with narrow therapeutic windows require higher confidence for the system to generate a dose recommendation. A routine antibiotic with a wide therapeutic range might require moderate confidence based on standard patient data. An anticoagulant with a narrow therapeutic window and severe adverse effects from dosing errors requires high confidence based on recent lab values, confirmed patient weight, and resolved drug interaction analysis.

The confidence computation integrates multiple inputs: data recency, measurement precision, inter-source agreement, drug interaction resolution completeness, and patient-specific risk factors. Each input contributes to the composite confidence state. A recent, precise lab value from a reliable source increases confidence. An old, estimated value from an uncertain source decreases it.

When composite confidence meets the risk-proportional threshold for the specific medication, the system generates the dose recommendation. When confidence falls below the threshold, the system enters non-executing mode for that recommendation: it explains what additional data would be needed to generate the recommendation with sufficient confidence, identifies which inputs are driving the uncertainty, and suggests clinical actions that would resolve the uncertainty.

## Non-executing mode in clinical context

Non-executing mode for pharmaceutical dosing does not mean the system stops functioning. It means the system stops generating specific dose recommendations and transitions to an advisory posture. Instead of recommending a specific dose, the system communicates what it knows, what it does not know, and what would resolve the uncertainty.

For the clinician, this transition is informative rather than disruptive. A system that says it cannot recommend a specific warfarin dose because the most recent INR is twelve hours old and the patient's hepatic function is uncertain provides actionable information. The clinician knows to order a current INR and assess hepatic function before proceeding. This is more useful than a specific dose recommendation generated from uncertain data that the clinician must independently assess for reliability.

The hysteretic recovery requirement means that once the system has entered non-executing mode for a high-risk medication, it does not immediately resume recommendations when a single data point is updated. The composite confidence must recover to a threshold above the pause threshold, ensuring that the underlying data quality has genuinely improved rather than marginally fluctuating around the boundary.

## Patient safety through structural governance

For healthcare systems deploying AI dosing tools, confidence governance provides the patient safety layer that responsible clinical AI requires. The system does not merely recommend doses. It governs its own recommendation authority based on its assessed confidence in the data underlying each recommendation.

Medication safety officers can configure risk-proportional thresholds based on institutional policies and pharmacological risk profiles. High-risk medications receive higher confidence thresholds. Specific patient populations, such as pediatric or geriatric patients with altered pharmacokinetics, receive adjusted thresholds. The governance is configurable and auditable.

For regulators evaluating clinical AI systems, confidence governance provides structural evidence that the system does not recommend doses when its confidence is insufficient for the clinical risk. Every recommendation carries its confidence state. Every non-executing transition is logged with the specific inputs that drove the confidence below threshold. The safety property is structural and auditable rather than dependent on clinician judgment to catch uncertain recommendations.

[Confidence Governance All 21 steps →](#)

Execution is a revocable permission, not a default.

Primary Technical Disclosure

[◦ Confidence-Governed Execution: When Agents Pause, Reassess, and Resume Safely](#)

Secondary Technical

[◦ Execution as Revocable Permission](#)[◦ Confidence as First-Class Computed State Variable](#)[◦ Composite Admissibility Evaluator](#)[◦ Confidence Trajectory Projection](#)[◦ Non-Executing Cognitive Mode](#)[◦ Task Class Differentiation Under Confidence Interruption](#)[◦ Confidence-Integrity Feedback Loop](#)[◦ Differential Rate Alarm Conditions](#)[◦ Hysteretic Authorization Recovery](#)[◦ Confidence Computation Function](#)[◦ Confidence-Driven Inquiry Mode](#)[◦ Curiosity as Confidence Modulator](#)[◦ Affect-Modulated Confidence Sensitivity](#)[◦ Effort Analysis and Path Optimization](#)[◦ Confidence-Modulated Discovery Traversal](#)[◦ Biological Signal to Confidence Coupling](#)[◦ Multi-Agent Confidence Propagation](#)[◦ Confidence-Governed Embodied Execution](#)[◦ Deferred Execution and Temporal Reauthorization](#)[◦ Execution Authorization Recovery](#)[◦ Confidence Contagion in Delegation](#)[◦ Confidence History Calibration](#)[◦ Attention Field](#)

Applications (General)

[◦ Autonomous Vehicle Execution Safety Through Confidence Gating](#)[◦ Clinical AI That Pauses When It Should Not Act](#)[◦ Confidence Governance for Nuclear Operations](#)[◦ Confidence Governance for Aviation Autopilot Systems](#)[● Confidence Governance for Pharmaceutical Dosing Systems](#)[◦ Confidence Governance for Bridge Structural Monitoring](#)[◦ Confidence Governance for Food Safety Inspection](#)[◦ Confidence Governance for Chemical Plant Operations](#)

Applications (Specific)

[◦ Agentforce Executes by Default](#)[◦ Microsoft Copilot Has No Confidence State](#)[◦ OpenAI Operator Cannot Govern Its Own Execution Authority](#)[◦ Claude's Safety Has No Computed Confidence Variable](#)[◦ Gemini's Multimodal Confidence Is Not Computed](#)[◦ Cohere Command Generates Without Computed Confidence](#)[◦ AWS Bedrock Guardrails Filter Output Without Governing Confidence](#)[◦ Azure Content Safety Classifies Harm Without Governing Execution](#)[◦ Google Vertex AI Safety Filters Without Confidence State](#)[◦ NVIDIA NeMo Guardrails Constrains Dialogue Without Governing Confidence](#)[◦ Guardrails AI Validates Output Without Governing Execution Authority](#)[◦ Lakera Guards Inputs Without Governing System Confidence](#)  
[Confidence Governance 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](mailto:nick@qu3ry.net)
- 72 28 14 36 01



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