

# The Incompatibility of Probabilistic Inference and Authority

Why AI Systems That Guess Cannot Be Trusted With Decisions

Timothy M. Gough  
IAMMOGO Intelligence Company, Inc.  
Fountain Valley, California, USA  
Email: [tgough@iammogo.com](mailto:tgough@iammogo.com)

**Abstract**—Modern artificial intelligence systems increasingly operate in roles with real-world consequences, yet most are built on probabilistic inference. This paper advances a central claim: probabilistic inference and execution authority are structurally incompatible. Systems designed to estimate likelihoods cannot reliably enforce permission, refusal, or fail-closed behavior, all of which are required for legitimate decision authority in high-consequence environments. Common approaches proposed to mitigate AI risk, including increased model scale, improved data quality, monitoring, audits, and explainability, do not resolve this incompatibility. These measures observe, analyze, or interpret outcomes after execution but do not govern whether execution should occur in the first place. As a result, they cannot prevent harm in systems where authority is exercised, only document it after the fact.

This paper argues that real-world AI systems require deterministic governance at the system level, independent of model behavior. Authority must be enforced before action occurs through explicit, rule-governed state transitions, such that denial results in no execution and no state transition. This work intentionally omits implementation detail and focuses on logical necessity rather than engineering approach, establishing a structural foundation for trustworthy AI in environments where accountability, safety, and control are non-negotiable.

## Table of Contents

- I. Introduction: The Authority Problem in Modern AI
- II. Terminology and Scope
- III. How Probabilistic Inference Works at a System Level
- IV. What Authority Requires in Decision Systems
- V. The Structural Conflict Between Inference and Authority
- VI. Logical Proof of Incompatibility

- VII. Why Scale, Tuning, and Explainability Cannot Fix This
- VIII. The Failure of Post-Hoc Governance
- IX. Deterministic Governance as a System Requirement
- X. Implications for AI Safety, Law, and Regulation
- XI. Conclusion: Authority Cannot Be Inferred
- References

## I. INTRODUCTION: THE AUTHORITY PROBLEM IN MODERN AI

Artificial intelligence systems were initially deployed as advisory tools. Early applications focused on recommendation, classification, and pattern detection, supporting human decision makers rather than replacing them. In these contexts, probabilistic inference was appropriate. Systems suggested possibilities, highlighted risks, and surfaced correlations, while humans retained authority over execution.

That boundary has eroded. AI systems are now routinely embedded in workflows where outputs directly trigger action. Automated systems determine hiring outcomes, credit approval, insurance eligibility, medical prioritization, content moderation, and access to public services. In these roles, AI no longer observes or recommends. It decides.

This shift introduces a fundamental problem of authority. Authority is not a measure of intelligence or accuracy. It is the power to permit or deny execution before action occurs. Any system exercising authority must be able to refuse action, fail safely under uncertainty, and enforce constraints regardless of output

confidence. Without these properties, authority cannot be meaningfully assigned or controlled.

Most modern AI systems lack enforceable permission boundaries. They are invoked to produce an output and are structurally optimized to do so. When uncertainty is high, they still respond. When inputs are incomplete, they still infer. When consequences are irreversible, they still execute. Responsibility is then distributed after the fact through monitoring, review, or explanation, rather than enforced at the moment of decision.

As AI systems continue to assume decision-making roles across critical domains, this absence of pre-execution governance creates systemic risk. The problem is not misuse or misalignment alone. It is architectural. Understanding why probabilistic inference cannot satisfy the requirements of authority is essential to defining what trustworthy AI must become.

## II. TERMINOLOGY AND SCOPE

Clear terminology is essential to avoid ambiguity when discussing authority, execution, and governance in intelligent systems. The following definitions are used consistently throughout this paper.

### AUTHORITY

Authority refers to the capacity of a system to permit or deny execution of an action before that action occurs. Authority is a control property, not a measure of intelligence, accuracy, or confidence. A system exercises authority only if it can enforce refusal, defer action under uncertainty, and fail in a closed state when required conditions are not met. A refusal expressed as an output does not constitute denial of execution if the system has already completed inference.

### PROBABILISTIC INFERENCE

Probabilistic inference is the process by which a system selects an output based on likelihood estimates derived from statistical patterns in data. Such systems are designed to produce an output when invoked, optimizing for selection rather than permission. Probabilistic inference may support prediction, classification, recommendation, or generation, but it does not inherently include execution control or refusal guarantees.

### DETERMINISM IN SYSTEM CONTEXT

Determinism, as used in this paper, does not refer to predicting outcomes or eliminating uncertainty about the external world. It refers to deterministic governance of system behavior. A deterministic system produces the same execution outcome for the same validated state and inputs, ensuring that permission to act is governed by explicit rules rather than inferred likelihood.

### EXECUTION VS OBSERVATION

Execution refers to the initiation of a state change that produces real world effect. Observation refers to analysis, monitoring, logging, explanation, or evaluation of behavior after execution has occurred. Systems that observe outcomes do not control execution. Governance based on observation is fundamentally distinct from governance that constrains action before it occurs.

### CLARIFICATION ON INDUSTRY PRACTICE

In practice, many production systems already separate probabilistic inference from execution authority. Probabilistic components are commonly used to estimate risk, likelihood, or preference, while deterministic logic is employed to determine whether inferred outputs are allowed to influence action. This separation reflects an implicit recognition that probabilistic inference alone cannot guarantee refusal, fail closed behavior, or reproducible execution.

This paper treats such designs as confirmation of the central claim rather than exceptions to it. In these systems, authority is not exercised by probabilistic inference itself, but by the deterministic structures that constrain when and how inference outputs may be acted upon.

### SCOPE CLARIFICATION

This paper examines the structural relationship between probabilistic inference and execution authority. It does not evaluate model performance, accuracy, or utility in advisory or exploratory roles. It does not propose specific implementations, algorithms, or enforcement mechanisms. The analysis is limited to logical necessity and system properties required for authority in high consequence decision environments.

The intent of this scope is to establish whether probabilistic inference can satisfy the requirements of authority by design, independent of scale, tuning, monitoring, or post execution oversight.

### III. How Probabilistic Inference Works at a System Level

Probabilistic inference systems operate by estimating likelihoods and selecting outputs based on statistical relationships observed in data. When such a system is invoked, its purpose is not to determine whether an action should occur, but to determine which output is most likely given the available input. The system is designed to complete a task by producing a result.

At the system level, this behavior is completion driven. Invocation triggers inference. Inference produces a ranked set of possibilities or a selected output. The process assumes that some output must be returned, even when uncertainty is high or information is incomplete. The system's success is measured by how well it selects among alternatives, not by whether it should act at all.

Likelihood selection is therefore the central organizing principle. Probabilistic systems evaluate inputs against learned patterns and choose outputs that maximize confidence or probability according to internal scoring criteria. This process does not encode permission, refusal, or prohibition as first class concepts. A low confidence output is still an output. Uncertainty influences ranking, not execution.

Crucially, probabilistic inference does not contain internal boundaries that prevent execution. The system does not possess a native concept of "not allowed." It cannot determine that producing an output is inappropriate in a given context. When refusal appears to occur, it is the result of selecting a refusal response as an output, not the result of blocking execution itself. The inference process has already completed.

This distinction matters because completion and control are different system properties. Completion answers the question of what output to produce. Control answers the question of whether output should be produced at all. Probabilistic inference is optimized for the former. It has no inherent mechanism to enforce the latter.

As probabilistic systems are embedded deeper into operational workflows, this completion driven behavior becomes increasingly consequential. When inference outputs are directly coupled to action, the absence of internal permission boundaries means that estimation is

allowed to propagate into execution by default. Probabilistic inference systems do not natively support a null execution state in which invocation results in no output and no state change. Any governance must therefore occur outside the inference process itself.

This structural characteristic is not a flaw or oversight. It is a defining feature of probabilistic inference. Understanding this behavior is essential to evaluating whether such systems can ever satisfy the requirements of execution authority.

### IV. What Authority Requires in Decision Systems

Authority in a decision system is not defined by intelligence, accuracy, or confidence. It is defined by control. A system exercises authority only if it possesses the power to permit or deny execution before an action occurs. Authority is therefore a pre-execution property, not a post-execution explanation.

For a system to hold real authority, refusal must be a first class capability. The system must be able to prevent action when required conditions are not met, when inputs are invalid or incomplete, or when uncertainty exceeds defined limits. This refusal cannot be optional or advisory. It must be enforceable, final, and independent of output confidence.

Fail-closed behavior is equally essential. In systems with real world impact, uncertainty must default to non-execution rather than action. A system that proceeds by default and attempts to correct or explain after the fact does not exercise authority. It exercises automation. Authority requires that the absence of permission results in no action.

Authority also requires consistency. Identical conditions must lead to identical execution outcomes. Without reproducibility, permission cannot be meaningfully audited or enforced. A system that produces different outcomes under the same conditions cannot guarantee that authority was applied rather than inferred.

These requirements are well established in safety-critical and regulated systems. Physical interlocks, transaction authorization systems, and access control mechanisms all operate on the same principle. Action is not allowed unless permission is explicitly

granted. Observation, explanation, or review after execution cannot substitute for this control.

When applied to intelligent systems, this distinction becomes decisive. A system that cannot refuse to act, cannot fail closed, and cannot guarantee consistent permission does not possess authority by design. It may assist decision making, but it cannot be trusted to govern execution.

Understanding what authority requires clarifies why probabilistic inference and decision authority diverge at a structural level. The next section examines this incompatibility directly.

## V. THE STRUCTURAL CONFLICT BETWEEN INFERENCE AND AUTHORITY

The conflict between probabilistic inference and authority arises from opposing system objectives. Inference systems are designed to select an output. Authority systems are designed to deny execution unless permission is granted. These goals are not merely different. They are structurally incompatible.

Probabilistic inference optimizes for selection. When invoked, the system evaluates inputs, estimates likelihoods, and chooses an output according to internal scoring criteria. Success is defined by the ability to return a result. Even low confidence, ambiguous, or uncertain conditions lead to selection rather than non-response. The system is organized around completion.

Authority requires the opposite capability. An authority-bearing system must be able to deny execution. It must block action when conditions are unmet, when information is insufficient, or when uncertainty exceeds acceptable bounds. Permission is not inferred. It is enforced. The absence of permission must reliably result in no action.

A system designed to always produce an output cannot simultaneously enforce permission. Selection assumes execution. Denial requires non-execution. Inference systems can choose a refusal response as an output, but that choice still occurs after inference has executed. The system has already acted by completing its task.

Authority, by contrast, requires the ability to prevent that execution entirely.

This distinction is not semantic. It is architectural. Selection driven systems answer the question of what to output. Authority driven systems answer the question of whether output is allowed at all. When inference is used to determine permission, permission becomes probabilistic rather than enforceable. The system no longer denies action. It estimates whether denial is likely.

As a result, any system that relies on probabilistic inference as the basis for execution authority will default to action under uncertainty. The structure of inference compels completion. The structure of authority demands restraint. These requirements cannot be satisfied by the same mechanism. Selecting a refusal response is still an act of execution; authority requires preventing execution from occurring at all.

This incompatibility explains why governance added around inference must remain external to it. It also explains why scale, tuning, or confidence thresholds cannot transform inference into authority. The next section formalizes this conflict as a logical necessity rather than an implementation concern.

## VI. LOGICAL PROOF OF INCOMPATIBILITY

The incompatibility between probabilistic inference and execution authority can be established through direct logical reasoning. This argument does not depend on system design choices, model size, or deployment context. It follows from the defining properties of inference and authority as system behaviors.

### PREMISE 1: PROBABILISTIC SYSTEMS MUST SELECT AN OUTPUT WHEN INVOKED

A probabilistic inference system is invoked for the purpose of producing an output. Its internal process evaluates inputs, ranks possible outcomes by likelihood, and selects a result. Completion is the objective. Even when uncertainty is high or confidence is low, the system is designed to return an output rather than to halt without resolution.

#### PREMISE 2: AUTHORITY REQUIRES THE ABILITY TO PRODUCE NO OUTPUT

A system that exercises authority must be able to deny execution. This denial must occur before action takes place and must result in no output, no execution, and no transition to a new system state. The absence of permission must reliably lead to non-execution. Without this capability, authority cannot be enforced.

#### PREMISE 3: SELECTION AND DENIAL ARE MUTUALLY EXCLUSIVE SYSTEM PROPERTIES

A system that must select an output in order to complete its function cannot guarantee non-execution. Even if the selected output is a refusal message, the system has still executed its core function. The inference process has occurred, and an output has been produced. Denial has been inferred, not enforced.

#### CONCLUSION: PROBABILISTIC INFERENCE CANNOT SATISFY AUTHORITY REQUIREMENTS

From these premises, it follows that a system whose operation requires output selection cannot simultaneously guarantee the absence of output required by authority. Probabilistic inference and execution authority therefore cannot coexist within the same decision mechanism. Authority requires the capacity to prevent execution altogether, while inference requires execution to complete.

This conclusion is a matter of logical necessity, not engineering limitation. No increase in scale, data quality, confidence calibration, or post-execution analysis can alter this relationship. As long as inference is invoked to determine execution, authority remains probabilistic rather than enforceable.

The implications of this incompatibility extend beyond system design. They determine whether intelligent systems can be trusted with decisions that carry real-world consequences. The next section examines why commonly proposed remedies fail to resolve this conflict.

#### VII. WHY SCALE, TUNING, AND EXPLAINABILITY CANNOT FIX THIS

In response to concerns about AI authority, several remedies are commonly proposed. These include increasing model scale, improving data quality,

introducing explainable AI techniques, and relying on monitoring or audits. While each may improve aspects of performance or transparency, none address the core incompatibility between probabilistic inference and execution authority.

#### LARGER MODELS

Increasing model size expands pattern coverage and output fluency, but it does not alter system behavior at a structural level. A larger probabilistic system still operates by selecting outputs based on likelihood. Scale refines estimation but does not introduce permission boundaries or refusal guarantees. A system that must select an output to complete its function remains unable to enforce non-execution, regardless of size.

#### BETTER DATA

Higher-quality or more representative data can improve predictive accuracy, but accuracy is not authority. Even a perfectly trained probabilistic system must still produce an output when invoked. Better data reduces error rates but does not grant the system the ability to deny execution when conditions are unmet or uncertainty is unacceptable. The requirement to select remains unchanged.

#### EXPLAINABLE AI

Explainability provides insight into why a system produced a particular output. It operates entirely after inference has occurred. Explanation does not prevent execution, nor does it introduce a mechanism for refusal. A system that can explain its decision after the fact still lacks the ability to enforce permission before action. Explanation clarifies outcomes but does not govern them.

#### MONITORING AND AUDITS

Monitoring, logging, and auditing observe system behavior after execution. These practices support accountability analysis, compliance review, and retrospective correction. They do not alter the conditions under which execution occurs. Observation cannot substitute for control. A system that must act before it can be reviewed cannot be said to enforce authority.

Each of these approaches improves visibility, performance, or oversight, but none introduce the defining requirement of authority: the ability to prevent

action entirely. They operate around inference rather than within execution control. As a result, they cannot resolve the structural conflict identified in this paper.

The persistence of these remedies reflects a misunderstanding of the problem. The issue is not insufficient intelligence, transparency, or supervision. It is the absence of enforceable permission at the moment of execution. Without that capability, no amount of scale, tuning, or explanation can transform probabilistic inference into authority.

The next section examines why governance applied after execution cannot compensate for this absence and why control must precede action.

### VIII. THE FAILURE OF POST-HOC GOVERNANCE

Post-hoc governance refers to any form of oversight applied after an action has already occurred. This includes logging, review, explanation, auditing, and retrospective analysis. While these mechanisms are valuable for understanding behavior and assigning responsibility, they do not constitute control. Observation after execution cannot prevent the execution itself.

Logging records what happened. Review evaluates whether it should have happened. Explanation attempts to justify why it happened. None of these alter the state change that has already occurred. In systems where decisions produce real-world effects, harm is not theoretical once execution completes. It is realized.

This distinction mirrors established principles in safety and law. Accident investigation does not prevent accidents. Financial audits do not stop unauthorized transactions. Legal appeals do not undo harm already inflicted. These processes exist to assess and respond, not to govern action at the moment it occurs.

In intelligent systems, post-hoc governance is often presented as a substitute for pre-execution control. Transparency, explainability, and review are treated as mechanisms of trust. However, a system that must act before it can be explained remains ungoverned at the critical moment. Authority is exercised without permission, and accountability is applied only after consequences are imposed.

Observation and control serve different functions. Observation answers the question of what occurred. Control answers the question of whether an action is allowed to occur at all. A system that lacks control cannot be made safe through observation alone, regardless of how detailed or comprehensive that observation may be.

For governance to be meaningful in decision systems, it must precede execution. Permission must be enforced before action, not inferred or justified afterward. Without this ordering, governance becomes reactive rather than preventative, and intelligent systems remain capable of causing harm even when their behavior is fully observable.

The next section turns to what is required to establish governance at the correct point in the decision process and why deterministic control is a necessary system property rather than an optional enhancement.

### 9. DETERMINISTIC GOVERNANCE AS A SYSTEM REQUIREMENT

The preceding sections establish that execution authority cannot be derived from probabilistic inference. If authority is to exist at all in intelligent systems, it must be enforced through a different system property. That property is deterministic governance.

Deterministic governance refers to the use of explicit, rule-bound control that governs whether execution is permitted before an action occurs. It is deterministic not because it predicts outcomes, but because it guarantees consistent enforcement. Under identical validated conditions, the same permission decision is reached every time. This consistency is what makes authority auditable, enforceable, and legally meaningful.

Prior work by IAMMOGO has articulated this distinction by separating intelligence from governance. In that framing, probabilistic models may generate insight, estimation, or recommendation, but they are not granted execution authority. Authority resides in an independent control layer that evaluates whether conditions for action are satisfied. This layer operates before execution and remains invariant to model confidence, output fluency, or statistical likelihood.

Deterministic governance must be independent of model behavior. A system that conditions permission on inference output collapses authority back into probability. Governance must instead evaluate state, input validity, and rule compliance directly. When requirements are unmet, execution is denied regardless of what the model produces. When requirements are met, execution is permitted through an explicit transition.

This ordering is critical. Enforcement before execution prevents harm rather than documenting it. Refusal becomes a guaranteed outcome rather than a probabilistic suggestion. Fail-closed behavior is enforced by design rather than approximated through thresholds or monitoring.

Deterministic governance does not eliminate uncertainty from intelligent systems. It constrains how uncertainty is allowed to influence action. Exploration remains possible. Recommendation remains useful. Authority, however, is exercised only when permission is unambiguous. Only deterministic permission decisions can be meaningfully audited, because probabilistic permission cannot be reproduced under identical conditions.

As intelligent systems continue to assume roles with real-world consequence, deterministic governance is no longer optional. It is a prerequisite for assigning responsibility, enforcing compliance, and establishing trust. Without it, authority remains inferred rather than controlled, and execution remains vulnerable to the structural limitations of probabilistic inference.

## X. IMPLICATIONS FOR AI SAFETY, LAW, AND REGULATION

As intelligent systems assume roles with real-world consequence, the absence of enforceable authority boundaries introduces growing legal and regulatory risk. Systems that act without pre-execution governance challenge existing frameworks for liability, accountability, and compliance, particularly in high-risk domains.

### LIABILITY

Liability depends on the ability to attribute

responsibility for an action. When execution is driven by probabilistic inference, responsibility becomes diffuse. Outcomes are attributed to models, data distributions, or system behavior rather than to an enforceable decision rule. This ambiguity complicates fault determination and weakens legal remedies. Without deterministic permission logic, it becomes unclear whether harm resulted from misuse, error, or a system acting as designed.

### ACCOUNTABILITY

Accountability requires that decisions be traceable to explicit authority. Post-hoc explanations and logs can describe what occurred, but they cannot demonstrate that an action was permitted to occur in the first place. In systems lacking deterministic governance, accountability is applied after execution, when consequences have already materialized. This reactive posture undermines institutional trust and shifts risk onto individuals affected by automated decisions.

### COMPLIANCE

Regulatory compliance increasingly demands evidence that systems operate within defined constraints. Probabilistic inference alone cannot guarantee consistent adherence to rules, thresholds, or prohibitions. Compliance mechanisms based on monitoring and reporting observe behavior after the fact, but they do not prevent unauthorized execution. Deterministic governance enables verifiable compliance by enforcing rules before action, producing outcomes that are reproducible and auditable by design.

### HIGH-RISK DOMAINS

In domains such as healthcare, finance, transportation, public services, and critical infrastructure, authority gaps carry disproportionate consequences. Decisions in these environments can affect safety, rights, and access to essential resources. Systems that default to action under uncertainty expose institutions to heightened legal exposure and systemic risk. As deployment expands, the cost of these authority gaps increases, both in human impact and regulatory response.

Across these areas, a common pattern emerges. Existing governance frameworks assume that authority is exercised through enforceable permission, not inferred likelihood. Intelligent systems that cannot satisfy this assumption strain legal definitions and regulatory

safeguards. As reliance on automated decision-making grows, the incompatibility between probabilistic inference and authority will become increasingly difficult to ignore.

Addressing this gap is not a matter of refining oversight or improving transparency alone. It requires systems capable of enforcing authority at the point of execution. The final section concludes by restating why authority cannot be inferred and why trust in AI depends on this distinction.

## XI. CONCLUSION: AUTHORITY CANNOT BE INFERRED

This paper has advanced a single, structural claim: probabilistic inference and execution authority are incompatible by design. Systems optimized to select outputs based on likelihood cannot guarantee the denial of execution required for legitimate authority. This incompatibility is not a matter of implementation quality, model scale, or data sufficiency. It is a consequence of how inference and authority function at a system level.

Authority is not an emergent property of intelligence. It cannot be inferred from confidence, accuracy, or explanation. Authority exists only where permission can be enforced before action occurs. A system that cannot refuse to act cannot be said to govern its actions, regardless of how well it predicts outcomes.

Efforts to address AI risk through scale, tuning, transparency, or post-hoc oversight do not resolve this limitation. These measures observe and interpret

behavior after execution. They do not control whether execution should occur. As intelligent systems increasingly participate in decisions with real-world consequence, this ordering becomes decisive.

Trust in AI does not depend on persuasion or performance alone. It depends on restraint. Systems that can stop themselves from acting when conditions are unmet reduce harm by design. Systems that cannot will continue to rely on inference where authority is required, leaving responsibility diffuse and risk unresolved.

The implication is unavoidable. If intelligent systems are to be trusted with decisions, they must be governed by mechanisms capable of enforcing refusal. Authority cannot be guessed, approximated, or explained after the fact. It must be enforced.

## REFERENCES

- Gough, T. M. (2025). *Beyond Probabilistic AI: A Deterministic Framework for Ethical, Explainable, and Fully Offline Machine Decision Systems*. Preprint, Version v1. Published November 30, 2025.
- Gough, T. (2025). *The Deterministic Unification Model: Completing AI Theory Through State-Transition Computation*. Report, Version 1.0. Published December 2, 2025.
- Gough, T. (2025). *A Universal Framework for Ethical Machine and Human Decision Systems: Deterministic Ethics-Constrained State Transition Law*. Preprint, Version 1.0. Published December 5, 2025.