

Perspective-Induced Openness: Why Embedded Agents Cannot Pre-Fix Their Own Futures

© 2025-2026 Michael Semprevivo

This work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0).
You are free to:

Share — copy and redistribute the material in any medium or format.
Adapt — remix, transform, and build upon the material for any purpose, even commercially.

Under the following terms:

Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

To view a copy of this license, visit: <http://creativecommons.org/licenses/by/4.0/>

Perspective-Induced Openness: Why Embedded Agents Cannot Pre-Fix Their Own Futures

Abstract

This paper isolates a source of indeterminacy that arises not from physical randomness or value underdetermination, but from the structural position of agents embedded within the systems they seek to predict. Framed within Informational Ontology (Rev5) and its derivative accounts of degeneracy and agency, this paper isolates a form of openness that follows from embedded self-reference and feedback rather than from indeterminism. While developed within this framework, the constraint identified applies to any system whose self-models participate in the dynamics they describe. Even in a globally deterministic universe, self-referential prediction and intervention generate genuine local openness for embedded agents. This openness is neither epistemic ignorance nor metaphysical indeterminism, but a necessary consequence of

self-location, feedback, and informational constraint. The account clarifies how determinism, deliberation, and agency coexist without contradiction.

1. Introduction: Openness Without Indeterminism

Debates about freedom, agency, and the openness of the future are often framed as disputes over determinism. If the universe is fully determined by prior physical states, it is commonly assumed that the future must already be fixed in every relevant sense. From this perspective, any appearance of openness is treated either as ignorance, illusion, or a residue of incomplete science. This paper rejects that framing. It argues that openness can be structurally real for agents, even in a globally deterministic universe, without invoking indeterminism, randomness, or metaphysical freedom.

The central claim is that openness arises from perspective, not from the laws of physics. Agents are not external observers of the systems they inhabit; they are embedded within them. As a result, agents cannot represent, predict, or pre-resolve their own future states in the same way an external description might. Attempts at self-prediction are not neutral observations but interventions that alter the informational and causal conditions of the system because for embedded systems, representation cannot be causally inert. This creates a form of local openness that is neither epistemic uncertainty nor ontological indeterminacy, but a necessary consequence of self-reference and feedback.

This paper presupposes the Informational Ontology (Rev5) and two derivative results already established elsewhere: first, that genuine decision-relevant states arise when internal constraints fail to uniquely specify a continuation (value degeneracy), and second, that agency consists in the resolution of such underdetermination through internally mediated constraint rather than executive choice. The present argument does not revisit those claims. Instead, it isolates a distinct explanatory layer: why, from within a system, certain futures cannot be pre-resolved even when the mechanisms of resolution are fully determined.

The target phenomenon is therefore not *how* decisions are resolved, but *why* they cannot be settled ahead of time from the agent's own perspective. The openness at issue persists even if every physical process unfolds deterministically, and even if an external observer could, in principle, describe the system's evolution completely. What matters is that an embedded system cannot occupy that external standpoint without altering the conditions being described. This limitation concerns pre-resolution from within and does not imply that the system's action space is underdetermined.

By separating global determinism from local perspective, this paper aims to clarify how deliberation, planning, and control remain non-illusory in a determined world. The future is not "open" because it is indeterminate in itself, but because it cannot be pre-resolved *from within* by systems that are part of what brings it about.

Framing: Relation to Other Derivative Papers

This paper occupies a distinct explanatory role within the Informational Ontology project. Whereas *Resolution Under Degeneracy* explains how underdetermined states are resolved through self-referential constraint closure, and *Agency, Salience, and Free Will* explains how such resolutions give rise to the appearance and practice of agency, the present paper addresses a different question: why such resolutions cannot be anticipated or pre-resolved *from within* by the agents whose actions bring them about. Perspective-induced openness is therefore neither a competing mechanism nor a phenomenological gloss, but an independent structural constraint arising from embedded self-reference.

2. Embedded Perspective and the Limits of Self-Prediction

Any account of openness that treats agents as external observers mischaracterizes the problem it seeks to explain. Agents are not positioned outside the systems whose futures they attempt to anticipate; they are components of those systems. This embeddedness places principled limits on self-prediction that do not arise from ignorance, computational weakness, or incomplete data, but from the structural conditions of self-location.

An external description of a system may, in principle, specify its future evolution completely. From such a standpoint, the system's state at a later time can be treated as already fixed relative to earlier conditions. However, an agent within the system cannot adopt this standpoint with respect to its own future without altering the system it describes. The act of prediction is itself a physical and informational event that feeds back into the system's ongoing dynamics. As a result, the future cannot be pre-resolved *from within* in the same way it may be described *from without*.

This limitation is not merely practical. Even a perfectly informed agent, equipped with complete knowledge of the relevant laws and initial conditions, cannot generate a representation of its own future behavior that is causally inert. Any such representation becomes part of the agent's internal state, influencing deliberation, action, and subsequent system evolution. The predicted future is therefore not independent of the prediction itself. The agent cannot treat its own future as a static object of observation without ceasing to function as an agent.

Self-prediction thus differs fundamentally from prediction of external systems. When an agent predicts the behavior of an object that does not depend on the prediction, the act of representation can, at least ideally, be separated from the causal chain being modeled. In contrast, when an agent predicts its own future actions, the prediction enters the causal loop it seeks to describe. This creates a structural asymmetry between third-person and first-person prediction that cannot be eliminated by additional information or computational power.

The consequence is a form of local openness rooted in perspective. The agent's future is not unsettled because the world lacks determinate structure, but because the agent cannot occupy a standpoint from which its own future is fixed independently of its present deliberation. From within the system, the future remains open in the precise sense that it cannot be pre-resolved *from within* without contradiction.

3. Self-Reference, Feedback, and Intervention

The limits of self-prediction described above arise not only from embedded perspective, but from the self-referential structure of agency itself. Agents do not merely represent possible futures; they act on those representations. As a result, representation and intervention are inseparable. Any forecast an agent forms about its own behavior functions simultaneously as a causal input into the process it seeks to anticipate.

This coupling between prediction and action introduces feedback. A representation of a future state alters the agent's present informational configuration, which in turn constrains deliberation and action. The system's evolution is therefore conditioned not only on external factors, but on its own anticipatory activity. In such systems, the future cannot be treated as an independent variable whose value can be fixed in advance without affecting the process of fixation itself.

Self-reference intensifies this effect. When an agent reasons about its own future behavior, the content of that reasoning becomes part of the causal story it is attempting to describe. Any internally generated prediction must be taken up by the same mechanisms responsible for selecting action. There is no clean separation between "model" and "modeled" in this case. The system refers to itself in a way that closes the loop between description and execution.

This feedback structure blocks pre-resolution *from within* even when the system's dynamics are fully determined. A deterministic rule governing state transitions does not license an internal representation of the future that leaves those transitions unchanged. The act of internalizing the rule and applying it to oneself changes the state to which the rule applies. Determinism at the level of laws therefore does not translate into fixity at the level of self-directed prediction.

Importantly, this is not a matter of instability, randomness, or noise. The interference arises precisely because the system is functioning according to its organizational role: representations guide action, and action reshapes future conditions. An agent capable of deliberation necessarily incorporates its own anticipatory states into its causal organization. This makes self-fixing futures impossible without disabling agency itself.

4. Perspective-Induced Openness as a Structural Condition

The forms of limitation described so far—embedded perspective and feedback through self-reference—together give rise to a distinct kind of openness. This openness is not reducible to ignorance, lack of information, or computational constraint, nor does it depend on indeterminism in the underlying laws. It is instead a structural condition imposed by the standpoint of systems that both represent and bring about their own futures.

Perspective-induced openness should be distinguished carefully from other sources of underdetermination. It is not value degeneracy, which concerns cases where internal constraints fail to uniquely specify an action. Nor is it epistemic uncertainty, which could in principle be reduced by additional information or improved inference. Even when internal constraints fully determine how resolution will occur, and even when all relevant information is available, the agent's future cannot be pre-resolved *from within* without altering the conditions of that resolution.

What defines perspective-induced openness is the impossibility of separating representation from participation. An agent cannot generate a complete, causally inert description of its own future trajectory because any such description becomes part of the system's present state. The future is therefore not an object that can be fixed independently of the processes that lead to it. This limitation is intrinsic to self-locating systems and does not disappear under idealization.

This kind of openness is local rather than global. From an external standpoint, the system's evolution may be entirely specified by prior conditions and governing rules. Nothing in perspective-induced openness requires denying that description. The claim is instead that the external description cannot be internalized wholesale by the agent it describes. The agent's standpoint imposes constraints that block the internal availability of a fixed future, even when such a future exists relative to an outside perspective.

Perspective-induced openness is also non-illusory. It is not merely a subjective feeling of freedom masking an already settled outcome. The openness corresponds to a real structural feature of the system: the agent's inability to pre-resolve its own future *from within* without self-interference. Deliberation is therefore not a post hoc rationalization of an outcome already fixed for the agent, but a genuine process through which futures are shaped under constraint.

5. Compatibility with Determinism and Informational Ontology

Perspective-induced openness does not conflict with determinism once the distinction between external description and internal standpoint is made explicit. Informational Ontology already

permits a fully determined global evolution while denying that such evolution must be available or actionable from every perspective. The present account sharpens that distinction by showing why internal availability of a fixed future is structurally blocked for embedded agents.

From the standpoint of Informational Ontology, determinism concerns the existence of lawful transitions between informational states. It does not entail that those transitions can be internally represented in advance by the systems undergoing them. A complete external description may specify how a system will evolve, but such a description is not equivalent to an internal state of that system. When internalized, descriptions become informational differences that participate in the system's dynamics rather than standing outside them.

This distinction resolves a common confusion. Determinism is often taken to imply that deliberation is redundant because the future is already fixed. That inference presupposes that fixity must be equally accessible from all perspectives. Perspective-induced openness shows why this is false. The fact that a future state is fixed relative to an external description does not mean it is fixed relative to the agent whose present actions help realize it. For the agent, the future remains open precisely because it cannot be pre-resolved *from within* without becoming causally entangled with the attempt to resolve it.

Within Informational Ontology, this openness complements rather than competes with earlier results. Value degeneracy explains why multiple continuations may be structurally available at certain internal states. Agency explains how such underdetermination is resolved through internally mediated constraint. Perspective-induced openness explains why that resolution cannot be anticipated in advance *from within*, even when the mechanism of resolution is itself determined.

6. Implications for Agency and Deliberation

If perspective-induced openness is a structural feature of embedded agents, then deliberation, planning, and control retain genuine significance even in a deterministic universe. These activities are not attempts to discover a future that is already fixed *for the agent*, but processes through which the agent participates in bringing about a future that cannot be pre-resolved *from within* its own standpoint.

Deliberation, on this account, is not rendered redundant by determinism. An agent's consideration of options, evaluation of consequences, and selection of actions are among the mechanisms through which constraint is applied over time. Because the agent cannot internally fix the outcome of these processes in advance without altering them, deliberation is neither idle nor illusory.

Planning similarly depends on perspective-induced openness. Plans function as provisional structures that guide action without determining outcomes independently of execution. A plan constrains without settling; it guides without foreclosing. The openness of the future is preserved

not despite planning, but because planning itself is a form of intervention that cannot be causally inert.

Control, finally, should be understood as the capacity to reliably influence outcomes across a range of possible futures, not as the ability to pre-fix a single trajectory. Agents exercise control by shaping how futures unfold, not by standing outside the process and selecting among already-fixed paths. Control is therefore compatible with determinism, but incompatible with total self-prediction.

By isolating perspective-induced openness, this paper shows how determinism, agency, and an open future can coexist without contradiction. Openness is not a defect in physical law, but a necessary feature of self-referential systems embedded in the worlds they help bring about.