Agency, Salience, and Free Will: A Constraint-Based Account

A Companion Paper to Informational Ontology (Rev 5)

Abstract

This paper develops a constraint-based account of agency, free will, and responsibility as downstream consequences of Informational Ontology (Rev 5). Free will is not treated as a faculty, substance, or metaphysical endowment, but as a structural condition that sometimes obtains when an agentive system confronts a future not fully specified by constraint. The account rejects libertarian indeterminism, hard determinism, homuncular models, and retributive desert while preserving a structurally non-illusory notion of choice where it occurs. Central to the framework is a distinction between constraint saturation, constraint bias, and underdetermination, as well as the role of salience as a pre-interpretive mechanism that can collapse the space of possible action prior to deliberation. Free will appears only under underdetermination—where multiple meaningful futures remain available, value structure is non-decisive, salience does not close the field of action, and sufficient temporal bandwidth exists for meaning to modulate transitions. Responsibility is shown to scale with available freedom rather than outcomes or metaphysical authorship. The resulting account reconciles phenomenology, moral practice, and structural rigor without revising the core ontology.

---

1. Introduction: Free Will Without Faculties

Debates about free will are often framed as disputes over metaphysical endowments. Either agents possess a special faculty—an uncaused power of choice that stands outside the causal order—or free will is dismissed as an illusion generated by complex but ultimately deterministic processes. Both positions assume that free will, if it exists at all, must be a general property of agents: something one either has or lacks.

This paper rejects that assumption as structurally unnecessary.

Within the framework of Informational Ontology (Rev 5), free will is not a substance, capacity, or inner faculty. It is not something agents possess continuously, nor something that must be metaphysically insulated from causation. Free will is instead a structural condition that sometimes obtains under specific configurations of constraint. When those conditions are absent, free will is not overridden, replaced, or commandeered by another mechanism—it is simply not instantiated.

The central claim defended here is that free will appears only under underdetermination, understood as the absence of constraint saturation. This does not mean that free will requires indifference, randomness, or symmetry between options. It requires that no combination of value, salience, urgency, or structural bias uniquely specifies a single future transition. Where the space of possible action remains open—where more than one meaningful trajectory is structurally reachable—willing becomes possible.

Constraint does not oppose agency by its mere presence. Constraint is what defines a space of possible action at all. Free will is not maximized by the absence of constraint, but eliminated by its saturation.

This account departs from both libertarian and compatibilist traditions. It does not posit a metaphysically exceptional chooser, nor does it reduce agency to post hoc rationalization of fixed outcomes. It avoids appeals to homunculi, hierarchical “selves,” or infinite regress. When constraint saturates the system, nothing takes over. The absence of free will is not a failure of control; it is a consequence of closure.

Crucially, this framework separates **free will** from **responsibility**. Responsibility does not require that free will be present at every moment of action. Agents may remain admirable, blameworthy, or answerable for actions that flow inevitably from their character, commitments, or history—even when no alternative was available at the moment of action. Free will grounds *maximal* responsibility, but responsibility itself scales with the freedom actually available, not with outcomes or metaphysical desert.

This paper develops a constraint-based account of agency and free will downstream of the canonical Informational Ontology chain:

Δ → R → I → A → V → M → P

No part of that ontology is revised or extended here. Agency, salience, free will, and culpability are treated as consequences of how constraint operates within awareness, value, meaning, and purpose.

Agency requires constraint to define a space of meaningful action, but free will appears only when that space is not fully closed. Constraint does not limit freedom by its mere presence; it limits freedom only when it saturates the space of possible action and forces a single outcome.

---

2. Constraint, Openness, and the Space of Action

Constraint is often treated as the enemy of freedom. On everyday and philosophical conceptions alike, freedom is implicitly equated with the removal of limits. Within Informational Ontology, this framing is inverted. Constraint is not what makes agency impossible; it is what makes agency intelligible.

A constraint, in the ontological sense used here, is any structural restriction on the space of possible transitions available to a system. Constraints define which continuations are possible, which are excluded, and how remaining possibilities are weighted. Without constraint, there is no differentiated possibility space—only undirected change. Such a state does not maximize freedom; it dissolves it.

Constraint performs three indispensable functions for agency:

1. **Differentiation** – Constraint carves a space of possible action by distinguishing one future from another. Without such differentiation, there are no alternatives to act between.

2. **Relevance** – Through value and meaning, constraint renders some futures non-equivalent to others. Action becomes intelligible as action *for reasons* rather than arbitrary motion.

3. **Persistence** – Constraint allows actions to propagate consequences across time. Without persistence, no transition can meaningfully count as an action.

What matters for free will is therefore not whether constraint exists, but **how it operates**.

2.1 Constraint Saturation, Bias, and Slack

Three structurally distinct regimes of constraint are relevant:

- **Constraint-saturated regimes** are those in which only one future trajectory is viable. Salience, urgency, coercion, or necessity closes the space of possible action such that

no alternative options exist.

- **Constraint-biased regimes** are those in which multiple trajectories remain possible, but strong asymmetries heavily favor one outcome. Alternatives exist but are costly or difficult to access.

- **Constraint-slack (underdetermined) regimes** are those in which no combination of value, salience, or structural bias uniquely determines a single transition. Multiple meaningful futures remain structurally reachable.

Free will appears only in the third regime. It is not the absence of constraint, but the absence of constraint **saturation**.

2.2 Underdetermination and Possible Action

Underdetermination should not be understood as indecision or symmetry. Preferences may exist, even strong ones, without forcing resolution. Underdetermination is a modal condition: the future is not fixed by the total constraint structure.

In simple systems, underdetermination may appear as a literal tie. In complex systems, it typically appears as a broader field of structurally reachable trajectories shaped by memory, expectation, social context, and meaning. What matters is whether more than one possible action remains available at the moment of transition.

---

3. Salience, Attention Capture, and Constraint Closure

Not all constraints operate through meaning. Some operate prior to interpretation by capturing attention and compressing time. Salience names this class of constraint.

Salience is the structural prominence of certain informational states within awareness. It elevates urgency, narrows attention, and shortens deliberative time. Salience is not a value judgment and not an act of choice. It operates by constraining which trajectories remain practically reachable.

Salience does not override free will; it prevents free will from arising by closing the space of possible action.

3.1 Salience and Meaning

Meaning organizes value across time. Through meaning, present states are related to past experience and anticipated futures. Meaning preserves context and optionality.

Salience contracts context. It demands immediacy and accelerates resolution. Hunger, pain, fear, panic, and addictive craving are paradigmatic not because they are irrational, but because they operate prior to interpretation.

When salience dominates, alternative trajectories may remain informationally represented, but they are no longer structurally reachable in action.

3.2 Closure Without Replacement

When salience closes the space of possible action, nothing replaces free will. There is no competing inner controller, impulse, or sub-agent. The system transitions because only one trajectory remains viable.

To suppose that some other mechanism must "take over" when free will fails is to mischaracterize the phenomenon. Closure explains action without remainder.

3.3 Temporal Compression

Salience operates in part by compressing time. Deliberation requires temporal bandwidth: the persistence of constraint slack long enough for meaning to influence transition. Salience shortens this bandwidth by forcing immediacy.

Salience may bias without fully closing the space of possible action, yielding diminished freedom. The degree of freedom scales with how much temporal and attentional space remains open.

---

4. Underdetermination and the Conditions for Free Will

Free will is a structural condition that obtains only when action is underdetermined by constraint.

4.1 Necessary Conditions

Four jointly necessary conditions characterize the presence of free will:

1. **Multiple Meaningful Futures**
   The system must face more than one future trajectory that is intelligible in terms of its values and meanings. Mere physical possibility is insufficient; the alternatives must be **differentially meaningful**—non-equivalent with respect to the system's value structure. Here, "meaningful" denotes differential relevance within the system's value structure, not subjective endorsement, judgment, or conscious evaluation.

2. **Non-Decisive Value Structure (Value Degeneracy)**
   Value differences must fail to uniquely force selection. Value degeneracy here names

non-decisiveness, not equality or indifference.

3. **Absence of Salience Closure**
   Salience must not close the space of possible action prior to interpretation.

4. **Sufficient Temporal Bandwidth**
   The constraint structure must remain open long enough for temporally extended value relations to influence the transition. Temporal bandwidth names a property of constraint persistence, not a psychological or computational resource.

If any condition fails, free will is not instantiated.

4.2 Degrees of Freedom

Because constraint operates by degree, free will admits of degree. Where bias is mild, freedom is robust. Where bias is strong, freedom is diminished. Where constraint saturates, freedom is absent.

4.3 No Executive Will

This account posits no executive self, higher-order will, or internal chooser. When the future is open, free will appears. When the future is closed, free will does not. Nothing else is required.

Accounts of free will often assume that genuine agency requires a selector: a rule, mechanism, or operation that determines which option is chosen among alternatives. This assumption is treated here as a category error. Free will, on this account, is not a selection process performed over options, but a structural condition under which more than one course of action remains available and the system nevertheless acts. When underdetermination obtains, action does not require further specification by a chooser, decision rule, or randomizer. The demand for such a mechanism merely reinstates the homunculus this framework explicitly denies, mistaking the absence of constraint saturation for an explanatory gap rather than recognizing it as the locus of agency itself.

When underdetermination obtains, no additional rule or randomizer specifies which option occurs. The resolution reflects the system's existing structure, history, and identity as a whole. No further explanation is required, because the demand for one presupposes a chooser beyond the system itself.

Underdetermination as used here is defined with respect to value decisiveness, not total causal ancestry; historical structure may constrain the space of possible action without uniquely specifying a single transition.

Nothing in this account denies that action under underdetermination is resolved through

self-referential informational constraints, as described in Informational Ontology (Rev 5); the present analysis concerns the conditions under which such resolution is free rather than forced.

---

5. Three Regimes of Action and Responsibility

Responsibility scales with the freedom structurally available at the moment of action.

- **Constraint-saturated action**: no alternative options exist; free will is absent; culpability is minimal.
- **Constraint-biased action**: alternatives exist under pressure; free will is diminished; responsibility is mitigated.
- **Underdetermined action**: multiple actions remain available; free will is present; responsibility is maximal.

This scaling avoids metaphysical desert while preserving accountability.

---

6. Phenomenology Without Illusion

Phenomenology tracks real differences in constraint structure.

- Compulsion corresponds to closure.
- Effort corresponds to bias.
- Freedom corresponds to openness.

Agents characteristically experience these regimes distinctly because the structure differs, not because freedom is illusory.

Ownership does not require ultimate origination. Actions may be owned without being freely chosen, and freely chosen without being metaphysically exceptional.

---

7. Boundary Cases and Applications

Addiction, coercion, training, and law are all explained by the same constraint-based machinery. Recovery restores freedom by reopening constraint slack, not by strengthening willpower. Coercion closes the space of possible action through salience amplification. Training reshapes constraint profiles, sometimes restoring freedom, sometimes replacing it with stable closure.

Artificial systems are neither granted nor denied free will here. The paper specifies conditions under which free will would obtain if such systems were to arise.

Nothing in this account requires that mechanisms of constraint operate identically across scales. Salience characterizes localized awareness; collective systems may exhibit closure through procedural, institutional, or structural mechanisms without phenomenological symmetry.

Conclusion

Free will is neither ubiquitous nor illusory where it occurs. It is local, fragile, and real where it occurs. It appears only where the future is open and disappears wherever constraint closes it. Responsibility scales with that openness, not with metaphysical authorship or desert.

If this account is correct, the long-standing question of whether free will “exists” has been misframed. Free will exists where the future is underdetermined—and only there.

Appendix A — Clarifying Common Objections

This appendix addresses several recurrent objections that arise when free will is assumed to require a selector, decision rule, or metaphysically privileged act of choosing. The purpose is not to defend the ontology adopted here, but to clarify how the present account departs from standard assumptions that often shape such objections.

A.1 “What Selects the Action?”

A common objection holds that any account of free will must specify a mechanism that selects one option rather than another. On this view, underdetermination demands a selector, whether deterministic, probabilistic, or random. This demand is rejected here as a category error.

Free will, on this account, is not a selection process performed over options. It is a structural condition under which more than one course of action remains available and the system nevertheless acts. When underdetermination obtains, no additional chooser, decision rule, or randomizer is required. The demand for one presupposes precisely the executive homunculus this framework denies, mistaking the absence of constraint saturation for an explanatory gap rather than recognizing it as the locus of agency itself.

A.2 "Is This Just Compatibilism?"

The present account is not compatibilist in the traditional sense. Classical compatibilism permits free action in fully determined scenarios, provided the action issues from the agent's internal states in the appropriate way. This framework explicitly denies that inevitability is compatible with free will.

On the constraint-based account developed here, free will is absent wherever the space of possible action is saturated, regardless of whether the determining factors are internal or external. Freedom requires underdetermination, not merely internal causation. While actions may still be attributable to the agent in constraint-saturated cases, they are not freely willed in the sense articulated in this paper.

A.3 "Is This Randomness or Noise?"

Another objection claims that if no decisive value gradient exists, any resolution must be arbitrary or random. This objection conflates situational openness with stochastic generation.

Randomness is a property of processes that generate outcomes by chance. Free will, as defined here, is a property of situations in which multiple actions remain structurally available. The account does not claim that agents randomly select among options. It claims that agents act under conditions where the future is not fully specified by constraint. No probabilistic sampling is implied or required.

A.4 "Does This Make Saints Less Free Than Moral Failures?"

On this account, agents whose character or training eliminates alternatives in particular situations may exercise less free will at those moments than agents who confront genuine openness. This consequence is acknowledged.

Free will is not identified with virtue, moral worth, or authorship of character. It is identified with local openness in the space of possible action. A highly virtuous agent may act inevitably in accordance with their character and still be admirable and responsible, even if free will is not instantiated at that moment. Conversely, free will may be present in cases of moral struggle without conferring virtue. These notions are intentionally separated.

A.5 “Is Underdetermination Merely Epistemic?”

The underdetermination at issue is not epistemic uncertainty or ignorance of hidden causes. It is structural non-closure. Salience, urgency, and temporal compression can genuinely eliminate alternatives, not merely obscure them. Conversely, where salience does not dominate and temporal bandwidth persists, multiple actions may remain structurally reachable even if the agent’s knowledge is imperfect.

Underdetermination here concerns the organization of constraint, not the limits of prediction.

End of Appendix A