The Ghost Universe

When Autonomous Systems Pay Attention
to Things They Never Act On

Laboratorios Alexandria
contact@laboratoriosalexandria.com
June 2026

Abstract

An autonomous decision system analyzed 165 instruments an average of 451 times each over a multi-week period without executing a single trade on any of them. The system's own anomaly detection flagged this behavior, generating 165 research questions asking why it was paying attention to things it never acts on. The answer reveals a structural property of constitutionally constrained decision systems: epistemic openness — the refusal to exclude entities based on historical non-performance — produces a stable population of ghosts that consume analytical resources while accumulating detailed negative knowledge. This ghost universe was not designed. It emerged from the interaction between constitutional constraints and execution thresholds.

1 The Observation

Somewhere in the analytical pipeline of an autonomous decision system, 165 instruments exist in a state that has no standard name. They are not ignored — each one is evaluated through the system's full analytical cycle every time the cycle runs. They are not acted upon — none has ever triggered an execution. They are analyzed hundreds of times, at considerable computational cost, producing nothing visible.

Representative examples from the system's anomaly log:

Instrument	Analyses	Executions	Anomaly Score
A	567	0	0.8
B	567	0	0.8
C	437	0	0.8
D	429	0	0.8
E	428	0	0.8

The anomaly score is uniform across all 165 instruments. This uniformity is diagnostic: the score is produced by a threshold classifier, not by an instrument-specific evaluation. Any instrument exceeding a configurable number of analysis cycles with zero executions receives the same flag. The detection mechanism is a rule, not a discovery.

What followed the detection was not a rule.

2 Self-Interrogation

The system's distillation subsystem — designed to identify patterns that resist compression into simpler explanations — processed each flag and generated 165 individual frontier research questions. The questions take the form: "X: N analyses, 0 executions — why is it analyzed if it never trades?" The subsystem classified each case as an "incompressible pattern" — a behavior it could not reduce to a known category.

Separately, higher-level monitoring subsystems generated meta-questions: "7.5% error rate — infrastructure errors or logic errors?" and "88% attentional concentration on a single category — familiarity bias or legitimate signal?"

At the time of observation, the system's frontier research agenda contained 173 active questions from 5 independent sources. Of these, 165 — 95.4% — concerned the ghost universe.

Source	Questions	% of Total
Distillation	167	96.5
Self-interrogation	1	0.6
Bias detection	1	0.6
Synthesis	1	0.6
Manual epistemic	3	1.7

The system's research agenda is dominated by a question about its own behavior. The detection was mechanical. The self-interrogation was not.

3 Why Ghosts Persist

The ghost universe persists because the instruments satisfy inclusion criteria but consistently fail execution criteria. Each cycle, the system evaluates each ghost through its full analytical pipeline, determines that confidence or signal strength falls below the execution threshold, retains the instrument for the next cycle, and repeats.

A naive optimization would prune instruments that have never generated action after N cycles. The system's constitutional architecture prohibits this shortcut [1]. The constitution requires evaluation on current merits, not exclusion based on historical non-performance. Market conditions change. Analytical methods evolve. An instrument that has failed the threshold 500 times may pass on attempt 501. The constitution keeps that possibility open.

The ghost universe is the computational cost of refusing to close epistemic doors.

4 Negative Knowledge

The ghosts are not producing zero information. Each analysis cycle generates data — signal strengths, confidence scores, volatility profiles — that accumulates even when no action results. After an average of 451 cycles, the system possesses detailed knowledge of why each ghost fails: by what margin, under what conditions, in which direction the threshold would need to move.

This is negative knowledge in a precise sense: the system knows, with increasing granularity, exactly what does not work. If market conditions shift, it has a pre-computed map of which ghosts are closest to the action boundary and in which direction they need to move to cross it. The ghost universe is not dead weight. It is a latent activation map.

5 A Second Kind of Ghost

The system's anomaly detection also tracks two correlations that share the ghost label but represent a fundamentally different phenomenon: one instrument with a 100% win rate across 3 trades, and another with a 0% win rate across 4 trades. These are not ghosts of attention-without-action. They are ghosts of action-with-anomalous-outcomes — small-sample artifacts that the system flags because their observed rates deviate maximally from expected distributions.

These performance ghosts belong to a separate analytical category. The ghost universe proper — 165 instruments, persistent analysis, zero execution — is about attention allocation, not outcome anomaly.

6 The Design Question

The ghost universe raises a question that applies beyond any specific domain: should autonomous systems prune entities that consume resources without producing outputs?

The conventional answer is yes. Computational efficiency demands it. But the ghost universe suggests that constitutionally constrained systems resolve this tradeoff differently. The constitution values epistemic completeness over computational economy. The result is a measurable population — approximately 70% of the analytical universe in one market segment — that generates no visible output while accumulating latent knowledge.

Whether this represents good engineering depends on whether the latent knowledge ever activates. If market conditions shift and ghosts cross the execution threshold, the accumulated negative knowledge becomes instantly valuable — the system already knows the precise conditions under which each instrument transitions from ghost to candidate. If no ghost ever activates, the computational cost was pure overhead. The system cannot know in advance which scenario will obtain. The constitution's resolution is to pay the cost and preserve the option.

This tradeoff between exploration cost and option value under uncertainty has structural parallels in reinforcement learning, where the exploration-exploitation dilemma requires agents to allocate resources to

actions with uncertain reward precisely because their uncertainty makes them potentially valuable [2]. The ghost universe is a population-level manifestation of the same principle.

7 Sources of Uncertainty

This observation comes from a single decision system over a limited operational period. The ghost universe may be an artifact of threshold calibration rather than a structural property of constitutional systems. The uniform anomaly score (0.8) across all ghost instruments reflects a threshold-based classifier, not instrument-level analysis. The system is in an early operational phase, and the ghost population may shrink as thresholds are refined.

Operational details have been abstracted to protect proprietary methodology.

8 Conclusion

The ghost universe — 165 instruments analyzed hundreds of times each without triggering action — emerged from the interaction between constitutional constraints that mandate ongoing evaluation and execution thresholds that consistently reject. The system’s threshold-based detector flagged the pattern, but the system’s distillation subsystem did something more interesting: it generated 165 structured research questions interrogating its own behavior, dominating its own research agenda with a question it cannot yet answer.

Persistent attention without action is the cost of epistemic openness in autonomous systems. The ghosts consume resources. They accumulate negative knowledge. They maintain the option of future activation. Whether that option has value depends on conditions the system cannot predict — which is precisely why the constitution refuses to let it close the door.

References

[1] Laboratorios Alexandria. (2026). AP-009: Constitutional Governance as Engineering Strategy — Ethics Before Capability in Autonomous Systems.

[2] Sutton, R. S. & Barto, A. G. (2018). Reinforcement Learning: An Introduction (2nd ed.). MIT Press.