

# Near-Knowledge

## What 434 Constitutional Deliberations Reveal About the Space Between Belief and Proof

Laboratorios Alexandria  
contact@laboratoriosalexandria.com  
June 2026

### ***Abstract***

Eight grade-B theses report confidence above 0.95. The deliberation system considers these propositions almost certainly true — and epistemically incomplete. This decoupling of confidence from completeness, observed across 434 thesis evaluations produced by a constitutionally governed multi-agent deliberation system, reveals a category of knowledge that existing frameworks do not adequately describe: propositions a system believes but cannot prove to its own standards. We call this near-knowledge. Of 434 evaluated theses, 85.0% received grade B (“plausible but insufficient for universality”) while 15.0% achieved grade A (“solid, verifiable”). The split is not explained by confidence alone — 19 grade-A theses (29.2%) have confidence below 0.80. The asymmetry emerges from the constitutional architecture itself: rules requiring verifiable evidence, explicit falsifiability, and declared uncertainty jointly create a deliberation funnel whose filtering properties can be studied independently of the domains under evaluation.

### **1 Introduction**

Eight propositions that the system considers 95% likely to be true did not earn the system’s highest grade. Nineteen propositions that the system considers less than 80% likely did.

These are not errors. They are the signature of a deliberation architecture that measures two distinct properties: how likely something is to be true (confidence) and how completely the truth can be demonstrated (epistemic grade). The gap between them — the space where belief outpaces proof — is where near-knowledge lives.

This paper reports empirical findings from a multi-agent deliberation system operating under a formal epistemic constitution. The system evaluates cross-domain scientific propositions through structured adversarial deliberation: domain specialists present evidence, a dialectician challenges assumptions, a synthesizer identifies structural patterns, and an epistemic judge grades the result. The constitution specifies what counts as a valid conclusion — verifiable evidence (Article 1), explicit falsifiability (Article 4), declared confidence with uncertainty sources (Article 7), and systemic rather than local optimization (Article 6). The constitutional governance approach and its engineering implications are described in AP-009 [1]. The system does not grade on a curve. It grades against a specification.

## 2 The Deliberation Funnel

At the time of analysis, the system had processed 2,376 deliberation sessions to completion. Of these, 1,943 (81.8%) were evaluated as grade C or D — the constitutional requirements for verifiable evidence and explicit falsifiability were not met — and produced no thesis. The remaining 433 sessions (18.2%) produced thesis evaluations graded A or B.

The grade distribution across 434 evaluated theses:

Grade	Count	%	Avg. Confidence	Range
A	65	15.0	0.801	0.40–1.00
B	369	85.0	0.673	0.30–1.00

The funnel has three levels. First, the upstream botanical pipeline filters correlations before they reach the Forum — of 5,375 correlations in the system, 92% have reached validated recurrence, and only those meeting maturity criteria are submitted for deliberation. The maturity gradient that governs this upstream attrition is documented in AP-010 [2]. Second, the Forum itself eliminates 81.8% of deliberated sessions at grades C and D. Third, of the surviving theses, 85.0% receive grade B rather than A.

The overall yield from processed deliberation to grade-A thesis is 65 out of 2,376 — 2.7%.

## 3 The Confidence Paradox

The relationship between confidence and grade is not what intuition predicts:

Confidence Band	Grade A	Grade B
0.95–1.00	25 (38.5%)	8 (2.2%)
0.90–0.94	9 (13.8%)	3 (0.8%)
0.85–0.89	11 (16.9%)	50 (13.6%)
0.80–0.84	1 (1.5%)	17 (4.6%)
Below 0.80	19 (29.2%)	291 (78.9%)

The overlap is substantial. Grade-A theses span the full confidence range — nearly a third report confidence below 0.80. Grade-B theses cluster at the bottom but eight reach the top. The system is not using confidence as a proxy for quality. It is measuring something else.

Examination of grade-B theses at high confidence reveals what that something else is. Recurring justifications for withholding grade A include: “solid in its functional scope but does not achieve the level of mathematical formalization required for an isomorphism”; “plausible but insufficient for universality”; “requires additional empirical validation.” These are not complaints about probability. They are complaints about completeness — the proposition may well be true, but its truth has not been demonstrated with the verification infrastructure the constitution demands.

The constitution creates a distinction between knowing and proving. A system that merely tracked likelihood would collapse confidence and grade into a single axis. This system does not. It maintains two axes because the constitution specifies requirements for proof that go beyond requirements for belief. This distinction echoes the difference between statistical significance and scientific validity that Ioannidis identified in the context of published research findings [3].

## 4 The Domain Effect

The funnel operates differently across domain pairs:

Domain Pair	Theses	A-Rate
Markets ↔ Computation & AI	18	33.3%
Life Sciences ↔ Validity of Conclusions	22	27.3%
Life Sciences ↔ Energy & Propulsion	44	25.0%
Life Sciences ↔ Space & Astrophysics	32	21.9%
Life Sciences ↔ Knowledge Frontiers	16	12.5%
Life Sciences ↔ Materials Science	184	10.3%
Life Sciences ↔ Epistemic Foundations	41	9.8%
Knowledge Frontiers ↔ Computation & AI	8	0.0%

The pattern is counterintuitive. Domain pairs with the highest A-rates are not the most familiar — they are the most distant. Markets and AI share little conceptual infrastructure; when they connect, the connection tends to be either clearly valid or clearly not. Life Sciences and Materials Science share extensive methodological overlap (nanofabrication, biosensors, drug delivery); the connections are numerous but harder to resolve definitively. Adjacent domains produce more gray zone. Distant domains produce more binary outcomes.

This finding is consistent with the atypical combination effect documented by Uzzi et al. [4], where scientific papers combining conventional and unusual domain pairings tend to produce higher-impact results — the novelty of distant combinations, when it survives quality filtering, carries stronger signal.

## 5 What Near-Knowledge Looks Like

Near-knowledge is not a theoretical construct. It has a specific empirical signature in the deliberation record.

A grade-B thesis at confidence 0.98 (THESIS-20260618-280) evaluating the relationship between Computation & AI and Materials Science reads: “solid in its functional scope, but does not achieve the level of mathematical formalization required for an isomorphism. The rarity of 0.99 suggests a robust physical correlation, but not a bijective map between domains.” The system believes the correlation is real. It cannot prove the correlation is structural rather than coincidental under the constitution’s standards.

A grade-B thesis at confidence 0.95 (THESIS-20260616-082) evaluating Market Regime and AI reads: “solid in the theoretical framework, based on ontological and systematic dynamic differences, but not supported by direct empirical evidence or formal mathematical models.” The system finds the argument theoretically convincing but lacks the empirical infrastructure to graduate it.

*Note: The quotations above are translated from Spanish-language system output. The originals read “sólida en su alcance funcional, pero no alcanza el nivel de formalización matemática requerida para un isomorfismo” and “sólida en el marco teórico, basada en diferencias ontológicas y dinámicas sistemáticas, pero no se sustenta en evidencia empírica directa ni en modelos matemáticos formales” respectively. They are direct translations of deliberation system output, not paraphrases.*

These are not rejected propositions. They are propositions in epistemic limbo — believed but unproven, preserved with their full uncertainty declarations, available for re-evaluation as evidence accumulates. The 369 grade-B theses constitute a structured inventory of what the system almost knows.

## 6 Sources of Uncertainty

The system has been operational for approximately one month. The 434 theses are a snapshot from an ongoing deliberation process. The constitutional constraints were designed by a single team and may encode biases not visible from within the system. The 85/15 ratio may shift as the knowledge base grows and deliberation agents encounter different domain structures.

The system’s current epistemic grade is F, reflecting its age. All findings are early empirical observations.

## 7 Conclusion

Constitutional governance in multi-agent deliberation creates a measurable gap between belief and proof. Of 434 evaluated theses, 85.0% fall into this gap — the system considers them plausible but cannot demonstrate their truth to constitutional standards. This gap is not a deficiency. It is the mechanism by which the constitution separates what a system believes from what it can show.

The deliberation funnel is steeper than the thesis distribution alone suggests. Of 2,376 fully processed sessions, 81.8% are eliminated at grades C and D before a thesis is generated. Of the 18.2% that produce theses, 85.0% receive grade B. The overall yield from deliberation to grade-A finding is 2.7%. This attrition rate is comparable to the maturity gradient observed in the upstream correlation pipeline [2] and to quality-producing funnels in drug discovery and venture capital — systems where the value of the output depends on the rigor of the filter.

The near-knowledge corpus — 369 theses with their confidence scores, falsifiability conditions, uncertainty declarations, and domain contexts — is the system’s structured map of what it almost knows. As evidence accumulates and the knowledge base deepens, some of these propositions will graduate to grade A. Others will be falsified. The constitution does not predict which. It ensures that the system knows the difference.

## References

- [1] Laboratorios Alexandria. (2026). AP-009: Constitutional Governance as Engineering Strategy — Ethics Before Capability in Autonomous Systems.
- [2] Laboratorios Alexandria. (2026). AP-010: The Maturity Gradient — Quantifying Epistemic Attrition in Cross-Domain Discovery.
- [3] Ioannidis, J. P. A. (2005). Why Most Published Research Findings Are False. *PLOS Medicine*, 2(8), e124.
- [4] Uzzi, B., Mukherjee, S., Stringer, M. & Jones, B. (2013). Atypical Combinations and Scientific Impact. *Science*, 342(6157), 468–472.

*Intellectual property for all systems, methods, and protocols described in this paper is proprietary to Laboratorios Alexandria. All rights reserved.*