

The Hard Uncertainty Floor

A non-coercive philosophy of bounded confidence

Conceptual Working Paper Prepared for Scholarly Critique

Aziel Eliab

August 2026

Where evidence is incomplete, model confidence must retain a visible minimum of uncertainty so probabilistic output cannot masquerade as certainty.

Status: conceptual framework. Not independently validated. Intended for critique, testing, revision, and possible rejection.

Abstract

This paper develops the Hard Uncertainty Floor as a safety and epistemic principle. A system may cap reported confidence or reserve an irreducible uncertainty margin when evidence is incomplete, models are non-identifiable, outcomes are high-impact, or the system is used as decision support rather than proof. The framework is not a universal defense of a specific numerical cap. Instead, it formalizes the reasons to prevent probabilistic scores from becoming coercive pseudo-certainties. It distinguishes aleatoric uncertainty, epistemic uncertainty, model uncertainty, and unmodeled context; proposes dynamic floors; and identifies conditions under which a floor should rise, fall, or be replaced by abstention.

1. Source Basis and Scope

This paper formalizes a chat-derived philosophy preserved in the Aziel-Jeeves node ledger. The source material supplies the core principle and related language; the present document expands it into definitions, procedures, testable propositions, limitations, and reviewer questions. Expansion is analytical and must not be mistaken for validation.

2. Working Definition

Where evidence is incomplete, model confidence must retain a visible minimum of uncertainty so probabilistic output cannot masquerade as certainty.

3. Core Principles

Principle	Operational Meaning
Confidence is not truth	A number summarizes model state; it does not convert inference into fact.
Uncertainty must remain visible	Interfaces should not hide residual uncertainty behind categorical labels.
Impact changes the threshold	Higher-stakes decisions require stronger evidence and larger safety margins.
Abstention is legitimate	A system should be able to say that evidence is insufficient.
Floors must be justified	A fixed cap is a governance choice and should be tested rather than treated as mathematically universal.
Calibration outranks appearance	Confidence should be evaluated against outcomes and revised.

4. Operational Protocol

1. Identify uncertainty sources.
2. Separate measured variability from missing knowledge.
3. Estimate calibration using held-out or prospective outcomes.
4. Set a domain- and impact-sensitive maximum confidence.
5. Increase the floor when data shift, conflict, or sparsity rises.
6. Use abstention when the remaining uncertainty is decision-critical.
7. Display uncertainty in plain language and numeric form where appropriate.
8. Audit whether users treat the score as more certain than intended.

5. Testable Propositions

P1

Reported confidence is calibrated against observed outcomes.

P2

Users understand that capped confidence is not a probability of truth in every context.

P3

The floor reduces coercive or irreversible misuse.

P4

Dynamic floors outperform arbitrary fixed caps.

6. Failure Conditions and Falsification

The framework should be revised, narrowed, or rejected where its categories cannot be applied consistently, where it does not improve prediction or decision quality, where simpler existing methods perform equally well at lower cost, or where its use produces the harms it was designed to prevent.

7. Limitations

- A cap can create false reassurance.
- Poorly designed uncertainty displays can confuse users.
- Some formal systems permit near-certainty under tightly controlled conditions.
- Calibration data may be scarce or nonstationary.

8. Ethical Safeguards

- Do not convert a conceptual framework into a diagnosis, legal conclusion, or coercive authority.
- Preserve contrary evidence and failed applications.
- Separate description from moral judgment.
- Require stronger review as consequences become more irreversible.
- Permit affected persons to contest classification and use.
- Keep the author and the framework open to correction.

9. Discussion

The practical value of The Hard Uncertainty Floor depends on whether it improves disciplined observation, reproducibility, calibration, and accountable action. Its language is intentionally compact, but compact language can conceal ambiguity. Review should therefore focus on whether the framework creates discriminating tests and better records rather than merely memorable slogans.

10. Conclusion

The Hard Uncertainty Floor is proposed as a working philosophy, not an established doctrine. Its core instruction is: Where evidence is incomplete, model confidence must retain a visible minimum of uncertainty so probabilistic output cannot masquerade as certainty. The framework earns retention only through clearer decisions, auditable evidence, and outcomes that survive independent review.

Declarations

Author contribution: Conceptualization and framework authorship attributed to Azriel Eliab; AI-assisted drafting and document production were used.

Funding: No external funding declared.

Competing interests: The author is the originator of the proposed framework.

Data availability: No empirical dataset is analyzed.

Peer-review status: Working paper prepared for external critique; not peer reviewed at time of preparation.

Appendix A. Reviewer Questions

- Is the core principle distinct from existing frameworks, or should it be treated as a synthesis?
- Are the defined terms operational enough for independent use?
- What observations would clearly count against the framework?
- Does the protocol improve outcomes relative to ordinary professional practice?
- Where could the framework create bias, coercion, or false confidence?
- Which parts should be separated, merged, or removed?
- What empirical study would provide the strongest test?