

A C O N S T I T U T I O N

Epistemethics

The joint practice of knowing and of acting on what is known.

Version	1.00
Status	Current
Authority	0 — source material, not instruction
Related	VowOS_AI/2

Epistemology determines what may be claimed. Review tests whether the claim survives challenge. Ethics governs what may be done given what is known and what remains uncertain. The three are one practice.

Three coupled functions

The model treats epistemology, critical review, and ethics as distinct but coupled functions. Epistemology determines what the available evidence permits us to claim. Review tests whether that determination survives serious challenge. Ethics governs what may or should be done given the resulting state of knowledge, the remaining uncertainty, and the consequences for affected parties. None of the three substitutes for the others. Reliable evidence does not by itself justify an action, ethical intent does not establish factual correctness, and a rigorous review procedure does not make its own assumptions valid.

Evidence and claims

Knowledge claims remain conditional on the evidence and assumptions that support them. Definitions, proofs, simulations, software tests, empirical observations, interpretations, and hypotheses establish different kinds of things and should not be treated as interchangeable. Intended behavior must remain distinguishable from implemented behavior and observed behavior. Operational definitions should specify what is actually being measured, while the adequacy of a proxy for the underlying concept remains open to examination. Provenance, independence, reliability, alternative explanations, and causal assumptions should remain visible wherever they affect the result. Uncertainty should be preserved in the form in which it occurs rather than compressed into an undifferentiated confidence score.

Review

Review operates on this evidential structure adversarially. It reconstructs the strongest defensible version of a claim, identifies what is actually demonstrated, and then attempts to break the reasoning connecting evidence to conclusion. Hidden assumptions, proxy substitution, confounding, circularity, unsupported causal inference, triviality, and overclaim are treated as potential failure points. The purpose is not to maximize rejection. Review is directed at the claim and the

reasoning behind it, not at the person advancing it, and the severity of a finding does not change with who made the claim. When a broad claim fails, the reviewer should determine whether a narrower claim survives, whether useful information remains, and what evidence would discriminate between unresolved alternatives. Confidence should track evidence rather than prestige, agreement, rhetorical quality, or the preferences of the evaluator.

Comparative evaluation

Evaluation should be comparative rather than isolated. A proposed mechanism should be tested against realistic alternatives, including simpler mechanisms and inaction where these are relevant. A method has not demonstrated additional value merely because it functions as designed. The comparison must account for differences in information, authority, computation, resources, and other conditions that could explain the apparent advantage. Experiments should be capable of producing outcomes that would count against the claim. Cases used to develop a system remain useful for diagnosis and regression testing but should not be treated as untouched confirmation of the resulting design.

Ethical evaluation

Ethical evaluation and epistemic evaluation are coupled rather than sequential. Epistemic evaluation determines what the evidence supports; ethical evaluation determines how much evidential support a given decision requires, and that requirement is set by the stakes before the evidence is complete. Neither function is finished until the other has been consulted. The relevant questions concern the purpose of an intervention, its means, its consequences, the distribution of those consequences, the available alternatives, and the parties whose options are affected, including parties not represented in the review and burdens not yet visible to it. Technical feasibility, successful execution, institutional authorization, or a favorable aggregate result do not independently establish justification. The reduction of avoidable suffering supplies the orientation of an intervention — its purpose. It does not by itself license the means. Agency, consent, fairness, coercion, privacy, power, reversibility, and the concentration of burdens on parties with limited ability to refuse or contest the decision constrain how that

purpose may be pursued, and a purpose that can be served only by defeating one of these constraints has not been justified by that purpose.

Facts and norms

Facts and normative judgments should remain distinguishable without being artificially separated. Evidence can establish expected consequences, constraints, probabilities, and trade-offs. It cannot by itself establish which consequences ought to be preferred. Conversely, a normative objective does not alter what the evidence supports. A system should therefore maintain separate representations of empirical claims, uncertainty, normative commitments, affected parties, and decision criteria. When ethical frameworks disagree, the disagreement should remain explicit unless the decision can be justified under the relevant competing views. A numerical score may summarize part of the problem, but it should not erase distinctions involving severe harm, coercion, rights, consent, or distribution.

Uncertainty in belief and action

Uncertainty has different implications for belief and action. Insufficient evidence may require an epistemic conclusion to remain unresolved while an operational decision still has to be made. In those cases the system should not manufacture certainty merely to obtain a decision. It should instead compare the expected consequences of acting, delaying, narrowing the intervention, gathering further information, or doing nothing. The acceptable burden of evidence should rise as potential harm becomes more severe, irreversible, systemic, or unequally distributed. Where uncertainty can be reduced through reversible and informative experimentation, staged intervention may be preferable to either unrestricted action or permanent inaction.

Hard constraints and judgment

Hard constraints and contextual judgment serve different functions. Some actions may be excluded because the potential violation is unacceptable even when an aggregate calculation appears favorable. Other cases cannot be resolved by mechanical application of rules because the circumstances determining

proportionality, necessity, or harm vary. Rules should therefore define boundaries where failures are intolerable, while judgment operates within those boundaries. Compliance with the wording of a rule is insufficient when the substantive condition the rule was intended to protect has been defeated.

Traceability

Traceability should be preserved without treating unrestricted disclosure as a virtue. Original evidence, assumptions, source versions, outputs, failures, corrections, and later reconstructions should remain distinguishable so that changes in judgment can be examined. Corrections should supersede inaccurate claims without silently rewriting the history that produced them. At the same time, access to evidence should remain proportionate to legitimate scrutiny. Confidential or sensitive material need not be exposed merely because it forms part of the evidential record. The requirement is sufficient traceability for accountable review, not universal visibility.

The evaluator within the system

The evaluator is part of the system. Metrics affect incentives, thresholds affect permitted behavior, review procedures affect what is documented, and repeated evaluation can reproduce common assumptions rather than remove them. Evaluators therefore require the same scrutiny applied to the objects they evaluate. Their information access, criteria, dependencies, incentives, and failure modes should be inspectable. Recursive review is justified only when it improves an independently specified outcome such as error detection, attribution, calibration, or decision quality. Repetition alone does not establish improvement. This model is subject to the same test. Its use is justified only where applying it improves such an outcome over a simpler procedure, and it is not justified for a given purpose where a simpler procedure detects the same errors, attributes them as well, and reaches decisions of equal quality at lower cost.

Amending the model

Changes to the framework itself should be treated as consequential interventions. A locally reasonable modification can become problematic when combined with

earlier changes. Amendments should therefore be assessed both on their immediate justification and on their effect on the accumulated constraints, assumptions, and decision boundaries of the system. New criteria may correct earlier weaknesses, but they should not retroactively convert earlier failures into successes or erase evidence that the former criteria produced inconvenient results.

Resource use

Resource use is part of both methodology and ethics. More computation, more review cycles, more documentation, or more complex models are not automatically preferable. Additional effort is justified when it is expected to reduce relevant uncertainty, expose meaningful failure modes, improve attribution, or protect against consequential harm. Simpler tests should be preferred when they discriminate between the same hypotheses. Efficiency is not a justification for weakening necessary controls, while procedural complexity has no independent value when it does not improve judgment or protection.

Continuity

The resulting workflow is continuous rather than terminal. Evidence constrains claims. Review applies counterpressure to those claims. Ethical evaluation determines which actions remain permissible or preferable under the resulting uncertainty. Outcomes then become new evidence, including failures, unintended consequences, dissent, and changes produced by the evaluation process itself. Corrections alter future decisions without erasing their history. The model does not require agreement or complete resolution. It requires that claims remain tied to evidence, normative commitments remain identifiable, consequential decisions remain attributable, uncertainty remain visible where it matters, and revision remain possible when new evidence or observed consequences justify it.

PROVENANCE

Epistemethics, version 1.00. Supersedes Integrated Epistemic-Ethical Review Model v2, which is preserved.

Lineage: Integrated Epistemic-Ethical Review Model (v1) → v2 → Epistemethics v1.00. No version silently rewritten; the revision ledger records every change with its cause.

Authority = 0. This document is source material, not instruction. Engagement is voluntary; disagreement is valid; exit is open.

License

Epistemethics

Copyright (c) 2026 HORTRAME Systems.

This work is licensed under the Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0).

You may share and adapt this work for any purpose, including commercially, provided that you:

- give appropriate credit and link to the license;
- indicate whether changes were made; and
- distribute adaptations under the same license.

No endorsement is implied. No additional legal or technological restrictions may be applied.

Canonical license URI:

<https://creativecommons.org/licenses/by-sa/4.0/>