

Consciousness Theory Evaluation Protocol (CTEP) Full Specification, Scoring Rubric, and Evaluator Guidance

Wharton, M. (2026). CTEP v1.0 Specification Repository. Zenodo. This document is the living specification for the Consciousness Theory Evaluation Protocol. It will be updated as the instrument develops in response to community application and critique. Version numbers are recorded below. The companion paper (Wharton 2026, Zenodo) should be consulted for the theoretical rationale behind each design decision.

Version history: v1.0 — June 2026 (initial deposit)

1. The Four Criteria

These criteria apply to every question in every tier. They are defined once and applied consistently.

Criterion	Definition	Key test
A: Coherence	The theory's answer must be internally consistent and consistent with its other commitments.	Does its answer contradict other questions?
B: Sufficiency	The answer must address the question rather than reframe, deflect, or change the subject.	Does the question change the subject?
C: Discriminability	The answer must generate at least one consequence — empirical or otherwise — that distinguishes it from alternatives.	Does the answer produce consequences that distinguish it from alternatives?
D: Parsimony	Assessed at two levels. Initial parsimony: ontological cost beyond what is required to answer the question. Secondary parsimony: explanatory power relative to other theories.	What does the theory cost, and how does it explain?

Dissolution Pathway Guidance (v1.0 provisional)

A dissolution response argues that the question is malformed rather than providing a direct answer. Three types are recognised. Type 1: the theory denies the explanandum and provides an account of why belief in it arises (e.g. strong illusionism). Pass if the replacement account is adequate. Type 2: the theory reframes the explanandum as a different phenomenon it does address. Evaluated on the adequacy of the reframed account. Type 3: the theory argues the question is malformed without providing a replacement. Fails sufficiency unless a formal argument for malformation is provided. Evaluators must specify which type of dissolution is being employed and assess accordingly. This guidance will be formalised as a scored sub-criterion in v1.1.

2. The Seven Questions and Tier Structure

Tier 1 — The Axioms: Load-bearing. Sufficiency failures here are not rescued by Tier 2 or 3 performance. Flag explicitly in synthesis.

Tier 2 — The Implications: Tests how well Tier 1 axioms map onto reality. Failures indicate refinement needed, not necessarily rejection.

Tier 3 — The Epistemics: Tests empirical tractability. Permanent failure renders a theory scientifically inert. Distinguish untestable in principle from untestable with current technology.

Q1: The Ontological Question (Tier 1)

Short form: What is consciousness?

What the theory must provide: A substantive account of what consciousness is: its nature, ontological status, and relationship to other fundamental categories. A purely functional definition is not a complete answer.

Cell-by-cell guidance:

Coherence: consistent with distribution, interface, phenomenal structure accounts?

Sufficiency: does it address what consciousness IS, not merely what it does or correlates with?

Discriminability: does it generate distinct consequences for downstream questions?

Parsimony: what primitives are introduced? What does the ontology dissolve or explain?

Q3: The Interface Question (Tier 1)

Short form: How does consciousness relate to physical processes?

What the theory must provide: A principled account of the relationship — identity, causation, correlation, constitution, filtration, or other — specifying direction and nature, not merely asserting a relationship exists.

Cell-by-cell guidance:

Coherence: consistent with ontological and distribution accounts?

Sufficiency: mechanism or structure specified, or merely labelled?

Discriminability: predictions about which systems are conscious, how variation maps to physical changes?

Parsimony: cost beyond background physics/biology? Downstream yield?

Q7: The Causal Function Question (Tier 1)

Short form: What does consciousness do?

What the theory must provide: An explicit position on causal efficacy. If causally efficacious: what work does it do that physical processes alone cannot? If epiphenomenal: why does it exist and why does it have the character it does? The philosophical zombie thought experiment is the canonical discriminability test.

Cell-by-cell guidance:

Coherence: consistent with interface and ontological accounts?

Sufficiency: specifies what consciousness does that unconscious processing cannot?

Discriminability: principled position on zombies; predictions about consciousness-absent conditions?

Parsimony: cost of positing the causal function? Downstream yield for agency and decision-making?

Q2: The Distribution Question (Tier 2)

Short form: What has consciousness?

What the theory must provide: Which systems are conscious, which are not, and why — derived from Tier 1 commitments. Must handle boundary cases: simple organisms, fetuses, vegetative state patients, artificial systems, split-brain patients, simple physical systems.

Cell-by-cell guidance:

Coherence: follows consistently from ontological and interface accounts?

Sufficiency: principled criterion applicable to novel cases?

Discriminability: predictions that differ from competitors for at least some systems?

Parsimony: initial cost of any domain expansions or restrictions? Downstream yield for evolution, development, pathology?

Q4: The Phenomenal Structure Question (Tier 2)

Short form: Why does consciousness have the character it does?

What the theory must provide: Two mandatory components. Component A — Qualitative character: why does the proposed mechanism produce THIS phenomenology rather than some other or none? Component B — Binding sub-question: how is the apparent unity of conscious experience explained? Specify whether binding account is mechanical, definitional, or eliminativist.

Cell-by-cell guidance:

Coherence: consistent with interface and distribution accounts?

Sufficiency: BOTH components addressed explicitly? Correlation without explanation fails Component A.

Discriminability: distinct predictions about phenomenal variation and binding breakdown conditions?

Parsimony: cost of explaining phenomenal structure and binding? Yield for variation, altered states, cross-species convergence?

Q5: The Cosmological Question (Tier 2)

Short form: What explains consciousness's presence in this universe at this time?

What the theory must provide: Account for consciousness in relation to cosmological and evolutionary history. The universe began ~13.8 billion years ago; biological complexity emerged only recently and locally. Explain consciousness's relationship to that history without contradiction.

Cell-by-cell guidance:

Coherence: consistent with ontological and distribution accounts? Generator theories must address the pre-biological universe. Filter theories must address what consciousness was prior to access architectures.

Sufficiency: principled explanation of phenomenological dimensions of cosmological history, not just functional complexity?

Discriminability: distinct predictions about when and where consciousness first appeared or was first accessed?

Parsimony: cost of any commitments about pre-biological or pre-cosmic status? Downstream yield?

Q6: The Empirical Question (Tier 3)

Short form: How could we tell?

What the theory must provide: Three mandatory components. A: Positive evidence — what observations or experiments would confirm or support the theory, specified precisely enough to distinguish confirmation from consistency? B: Falsification conditions — what would falsify the theory? Failure to specify is a component B fail. C: Counter-evidence handling — how does the theory account for unconscious processing evidence, split-brain findings, vegetative state cases, and cross-species anomalies?

Cell-by-cell guidance:

Coherence: predictions consistent with Tier 1 and 2 commitments?

Sufficiency: predictions test the theory's distinctive claims, not claims shared by all theories?

Discriminability: predictions differ from at least one significant competitor?

Parsimony: infrastructure cost of testing? Downstream yield — how many empirical questions do the predictions address simultaneously?

3. Scoring Framework

Each of the seven questions is assessed against each of the four criteria, producing twenty-eight evaluation cells. Parsimony is reported qualitatively at two levels rather than scored Pass/Partial/Fail.

Score	Meaning	Evaluator requirement
Pass	Substantive, adequate response	Written justification required. State what the response achieves.
Partial	Genuine but incomplete response	Written justification required. State what the response achieves AND what it leaves unresolved
Fail	Inadequate response	Written justification required. Specify failure type: omission (theory does not address), evasion

4. The Synthesis

After completing all twenty-eight cells, the evaluator produces a structured synthesis addressing four questions in order.

Tier 1 Integrity: Do the three Tier 1 questions cohere with each other? A theory whose ontological, interface, and causal function accounts are mutually consistent has Tier 1 integrity even if individual cells contain weaknesses. Flag contradictions explicitly.

Characteristic Profile: Where is the theory's explanatory power concentrated and where are its liabilities? Identify the strongest two or three cells and weakest two or three cells. Explain what the pattern reveals about the theory's underlying commitments.

Comparative Standing: How does the theory's profile compare to its closest competitors? Which competitor theories share the same strengths and weaknesses? Which cells most clearly differentiate this theory from those competitors?

Research Frontier: What would it take to move the theory's Partial scores to Pass scores? Identify the most tractable open questions where targeted theoretical or empirical work could materially strengthen the theory's standing on the protocol.

5. Versioning and Attribution Requirements

Every application of CTEP must record the following before the evaluation begins:

- Theory being evaluated and specific version or source text used
- Date of evaluation
- Evaluator's declared theoretical commitments and disciplinary background
- Evaluator's prior exposure to this theory or its proponents
- Version of the CTEP instrument applied (this document)
- For LLM evaluators: model identifier, version, training data cutoff date, any relevant system prompt conditions

Declarations are not disqualification criteria. They are transparency mechanisms. An evaluator with strong prior commitments may still produce a valid profile; their priors are recorded so that the synthesis layer can identify prior-correlated divergences as contested findings rather than instrument errors.

6. Parsimony Downstream Yield Heuristic (v1.0 provisional)

Downstream yield is measured by counting the number of distinct explanatory problems the commitment dissolves or addresses that no lower-cost commitment addresses. The key qualifier is "that no lower-cost commitment addresses" — this prevents a theory from claiming downstream yield for explananda that any theory handles. A problem counts as a distinct unit of downstream yield if: (a) it is recognised as a genuine explanatory demand in the literature, (b) the theory provides a principled account of it rather than merely predicting its occurrence, and (c) the account derives from the specific commitment under assessment rather than from background commitments shared by competitors. This heuristic will be formalised as a structured checklist in v1.1.