# Binding Engines and the Absurd: A Preregistered, Multimodel Pilot Study of Faith Traditions as Mythologies

Scott W. Waddell, D.S.I.

August 23, 2026

## Author Note

Waddell is the sole author and held sole human decision authority for the study. Seven model variants from six families provided role-separated AI research assistance. Their full governed identities and the short forms used thereafter are: Anthropic claude-fable-5 (**Fable 5**); Anthropic claude-opus-5 (**Opus 5**); Zhipu/GLM **GLM 5.2** (glm-5.2; **GLM**); NVIDIA **Nemotron 3 Ultra** (nemotron-3-ultra; **Nemotron**); DeepSeek deepseek-v4-pro (**DeepSeek**); Moonshot AI kimi-k3 (**Kimi**); and OpenAI gpt-5.6-sol (**gpt-5.6**). The audit record retains their recorded identifiers. Model assistance does not confer byline authorship or decision authority; attributed model artifacts and judgments remain identified in the Method section and audit record.

The OpenAI assistant that produced this revision also drafted the paper's account of model-role selection. Its prior roles include source verification, existential-judgment drafting, paper review and critique, figure production, sixth-revision finalization, and R7 production and lead QA. It therefore reconstructed the account, including its own assignments, from the governed decision ledger rather than from memory or self-assessment. Its R8 lead QA was author-side only; Kimi held the reserved exact-hash cross-family review gate, and that review was completed before acceptance.

Waddell is a member of The Church of Jesus Christ of Latter-day Saints. He reports greater familiarity with LDS doctrine and official sources than with the doctrines and source hierarchies of the other religious cases. That self-reported asymmetry is both an interpretive

resource and a limitation: it may make the LDS source hierarchy more familiar to him while increasing the risk that omissions elsewhere go unrecognized. Role-separated canon review, coding, adjudication, and targeted source verification reduced dependence on his interpretation after the design choices had been made. They did not control his choice of research questions, case set, dimensions, framework, or initial source architecture; make the models neutral; or establish equivalent corpus completeness or measurement across traditions.

**Abstract**

Faith traditions join postmortem claims that cannot be tested during life to prescribed practices and representations with observable premortem consequences. This pilot asks whether their structural content differentiates social binding under a registered practice-comparability rule and how canonical texts and prescribed practices structure Camus's absurd confrontation. Nine cases, five outcome-eligible religious traditions and four registered comparators, were coded under a calibrated instrument by two AI coders blind to the calibration key and to each other's submissions, adjudicated cross-family, and analyzed under preregistered rules. Seven model variants from six families occupied disclosed, role-separated research-assistance positions; Waddell retained sole authorship and decision authority. On the social axis, a crisp-set QCA using US tradition-level retention as the binding outcome produced three one-case configurations with BIND = 1 and one contradictory all-zero configuration. No Boolean minimization was possible, no row is assertable as a finding, and excluding Theravāda, the weakest outcome proxy and a practice-matching failure, removed all outcome variance. $H_0$-social is not rejected: the study provides no evidence against it, and content-based binding found no support at N = 5. On the existential axis, three religious cases were classified LEAP, two were UNDECIDABLE-ON-RECORD, and none received SUSTAINS; only the registered Camusian comparator received

SUSTAINS. $H_0$-existential is rejected (argued) only on its universal reading, solely through that comparator, while it is not rejected on the religious-mythology reading; the comparator-carried rejection is "tautology-adjacent" as evidence about religion, and the practice-bridge interpretation remains contestable. No joint analysis was performed, so RQ3 remains untested. The methodological contribution is a provenance design auditable within the governed workspace, making role exposure, disagreement, and repair history inspectable. It demonstrates traceability and failure visibility, not empirical validation or superiority to conventional research production.

## 1. Introduction

Comparing faith traditions empirically presents a partition problem. Their postmortem truth claims cannot be tested during life, but their prescribed practices and public representations may have observable premortem consequences. Treating those domains as one evidentiary object either asks social evidence to adjudicate metaphysics or lets untestable claims determine an empirical comparison. This study therefore brackets postmortem truth and registers three questions: whether structural features of mythological content covary with binding when practice intensity is held comparable; how canonical texts and prescribed practices structure the confrontation between the human demand for meaning and the world's silence; and how binding and lucidity relate. The present record analyzes the first two. RQ3, the joint question, remains untested.

The motivating tension is deliberate, and it is the author's own. On one side stands a content prediction developed in earlier work on Latter-day Saint mythology read through Durkheim: a corpus of stratified heavens, binding covenants, and family bonds sealed past death should bind its adherents more strongly than doctrinally thinner alternatives. On the other side

stands an evidential discipline argued in earlier work on death as the final empirical test: postmortem hypotheses are unfalsifiable during life, so traditions should be evaluated through their premortem, observable life-hypotheses. The research design formalized that collision by treating faith traditions as mythologies in the neutral, Durkheimian sense, collective representation systems with prescribed practice, and by crossing two hostile nulls: a strict Durkheimian null on the social axis and a strict Camusian null on the existential axis. Postmortem claims enter only as textual structures; their truth remains bracketed.

This is a pilot of that design: nine cases, five with an observable binding outcome; one wave of one US survey; and a deliberately minimal analytic ceiling registered before analysis. The scale permits a disciplined pilot, not a general claim about religion's social or existential effects.

The paper makes two bounded contributions. Substantively, it reports no evidence against $H_0$-social and no support for content-based binding at $N = 5$; on the existential axis, it rejects $H_0$-existential (argued) only on the universal reading, solely through the Camus comparator, while not rejecting it on the religious-mythology reading. No joint analysis was performed, so RQ3 remains untested. Methodologically, it documents a role-separated multimodel provenance apparatus that makes exposures, disagreements, and repairs inspectable. That apparatus establishes traceability and failure visibility, not empirical validation or superiority to conventional research production.

## 2. Theoretical framework: two reductions, crossed

The design's engine is a deliberate crossing of two thinkers who examined religion naturalistically from opposite ends. Durkheim asked what religion *does* while bracketing what it *claims*; Camus asked what its claims *cost* while ignoring what its gatherings produce. The design

installs each thinker's own reduction as the null hypothesis on his own axis. This section integrates the load-bearing debates behind that choice directly.

### 2.1 Durkheim's wager: the group binding itself

Durkheim's sociology of religion rests on a single audacious identification. In *The Elementary Forms of Religious Life* (Durkheim [1912] 1995), the division of the world into sacred and profane is the elementary religious fact; the experience that generates and renews it is **collective effervescence**, the heightened state of assembled, synchronized practice; and the sacred object is, in the end, the group representing itself to itself. Religion is real, on this account, because society is real: the force worshippers feel is the genuine force of their own collectivity, symbolized. Two earlier works supply the outcome side of the design. From *Suicide* (Durkheim [1897] 1951) come **integration** and **regulation** as separable, observable binding variables, with anomie as the diagnostic failure state; from *The Division of Labor in Society* (Durkheim [1893] 1984) comes the mechanical/organic solidarity typology. The codebook's Durkheimian dimensions descend directly: sacred/profane boundary and access architecture (D1), effervescence machinery (D2), integration (D3), regulation and anomie response (D4), solidarity type (D5), totemic anchor (D6).

#### *Plain-language guide: Durkheim and "society as God"*

The shorthand "society as God" does not mean that a community is literally a deity. That is not how this study uses the idea. Durkheim's point is about social power. A community is larger than any one member. It gives people shared rules, duties, stories, and a sense of who they are. When people gather and act together, that power can feel sacred. Religious stories and

symbols give the group's moral life a visible form. In this limited sense, people can experience the real force of their community through the sacred things they honor.

This does not mean that believers knowingly worship themselves. It also does not prove or disprove the existence of God. This study sets that question aside. It uses Durkheim's framework to ask how shared stories, rituals, symbols, and rules help hold a group together.

The key terms have narrow meanings here:

- **Sacred** means set apart and protected, not simply good. **Profane** means ordinary, not bad. A sacred/profane boundary is the line a group draws between them.
- **Collective effervescence** is the strong shared feeling that can arise when people gather and act together.
- **Integration** asks how firmly people are tied to the group. **Regulation** asks how strongly shared rules guide them. **Anomie** is a loss of clear rules or social direction.

The study tests a deliberately strict version of this view: when levels of shared practice are similar, the particular religious story may not explain differences in social binding. That is a starting claim to test. It is not a finding, and it is not the whole of Durkheim's theory.

Three successor literatures sharpen the wager into measurable form. Collins (2004) turned effervescence from description into mechanism: interaction ritual chains with specified inputs (copresence, mutual focus, shared mood, barriers to outsiders) and outputs (emotional energy, sacred objects, moral boundaries). D2's record format is Collins's input list operationalized for texts. The neo-Durkheimian extension of the sacred beyond congregational religion (Bellah 1967 on American civil religion; Alexander and Smith 2005) licenses the secular-civic-ritual comparator: if civil ceremony can run the effervescence engine on nontranscendent content, it is exactly the practice-without-transcendent-content cell the factorial

contrast needs. And Berger (1967) bridges the two axes in advance: a religious *nomos* shields against chaos and depends on a social plausibility structure, predicting the D4/C3 correlation (strong anomie response co-occurring with deflection of inquiry) that the joint framework treats as an open question rather than an assumption. Lukes (1985) supplies the standing caution the design institutionalizes as its first named threat: Durkheimian functionalism risks circularity, which is why mythology is defined structurally and binding measured strictly as an outcome.

Read strictly, Durkheim's identification is deflationary about content: if the referent of religious symbolism is the group itself, then the *specific* story told, whether stratified or binary eschatology, covenant or karma, is substitutable, and only the telling-together does the binding. The study is explicit that installing this strict reading as the null, and treating content and practice as separable at all, is **a defended design decision, not an inherited default**, with the burden-of-proof asymmetry it creates acknowledged rather than hidden.

### 2.2 Camus's verdict: the absurd and its exits

Camus's starting point is not a theory of religion but a relation. The **absurd** is neither the human demand for meaning nor the world's unreasonable silence taken alone; it is their confrontation, and it exists only while both terms are held (Camus 1991a). *The Myth of Sisyphus* then audits the exits: suicide, the leap, and revolt. What Camus calls **philosophical suicide** abolishes one term of the confrontation by asserting that the silence answers, whether with a God, a meaning, or a beyond, supplied by fiat. Its purchase price is named precisely: **lucidity is traded for hope**. The residual option, **revolt**, keeps both terms alive, refusing suicide and hope alike, sustaining the confrontation without pretending it has been answered. *The Rebel* extends revolt into a social thesis of shared being (Camus 1991b), a solidarity grounded in shared refusal rather than shared sacred apparatus, which the design registers as $H_2$-social and codes as a

solidarity source (C5). The study's reading of these as *strategies differing in structure*, not merely in mood, follows the main line of Camus scholarship (Sagi 2002 on the absurd as a lived tension different life-forms sustain or collapse; Foley 2008 on the internal argument from absurd to revolt).

***Plain-language guide: Camusian absurdism***

Absurdism begins with a clash. People want life to have a clear, final meaning, but the world gives no final answer that can be proved. Camus calls this clash **the absurd**. He does not mean that life is silly, that nothing matters, or that people should give up.

**Lucidity** means seeing the clash clearly. A **lucidity cost** asks whether close questioning must be limited or turned away for an answer to hold. A **leap** closes the clash by claiming a final answer that the world's silence has not proved. Camus calls this move **philosophical suicide**. He does not mean physical suicide. **Revolt** means refusing both self-destruction and an unproved final answer. It does not require violence. A person keeps living, choosing, and caring while admitting that the final answer remains uncertain. In this framework, **hope** has a narrow meaning: a final answer used to close the clash, not every ordinary wish or plan.

This study uses these ideas to sort how canonical texts and prescribed practices handle the clash. It asks whether they keep the question open, close it, or leave the answer unclear. It does not decide whether a religion is true or false, and it does not claim to know what individual believers privately think or feel.

The existential axis had a rival frame, and the record says so: terror management theory (Becker 1973; Solomon et al. 2015) is the empirically dominant mortality-and-meaning literature. The design declined it for two stated reasons. TMT explains *why* worldviews are held (anxiety buffering), not *how a given corpus structures the confrontation*, and only the latter is

codeable from canonical texts under the document-only constraint. And TMT's outcome variable is the private state the representation stipulation rules out of evidence. TMT enters as background theory predicting that some resolution apparatus will be present in every durable tradition, a prediction compatible with both nulls, while the Camusian typology supplies the codeable distinctions among apparatuses. The codebook's existential categories descend accordingly: absurd acknowledgment (C1), resolution type (C2, descriptive only), lucidity cost (C3), status of hope (C4).

**2.3 The crossing: each reduction as the other axis's null**

The two reductions are made to police each other. On the social axis, *Durkheim's* reduction, that content is interchangeable and practice binds, is the null, and content-differentiation must earn rejection; on the existential axis, *Camus's* reduction, that every religious resolution is the same leap, is the null, and kind-difference must earn rejection. Each thinker also arms an alternative against the other's null: Camus's revolt-solidarity stands as $H_2$-social against Durkheim's sacred apparatus, and the Durkheimian possibility of ritual running at full power without doctrinal assent stands behind $H_2$-existential (via the representation discipline of Chaves 2010 and the ritual-over-sincerity literature). The study discloses that this symmetry is structural-role symmetry, not minimal-assumption symmetry: the Durkheimian null is the more parsimonious claim for this unit of analysis, and the asymmetry is carried openly.

Neither null is a straw man; each is one pole of a live research program. The practice pole has a bibliography: costly-signaling theory (Sosis and Alcorta 2003) makes participation expense the binder and belief content the badge; Whitehouse's (2004) modes of religiosity make the form and cadence of ritual, doctrinal versus imagistic transmission, predict social morphology with propositional content epiphenomenal. Read strictly, both *are* $H_0$-social. The content pole

answers: rational-choice theorists (Stark and Finke 2000) make the specific compensators a tradition offers load-bearing by construction, and Norenzayan (2013) argues that one content feature, morally concerned and watchful "Big Gods," scaled cooperation itself: a covariation claim between a structural content feature and a binding outcome, which is $H_1$-social argued at civilizational scale. The cognitive science of religion sits between (Boyer 2001; Atran 2002): content is constrained by evolved cognition, which explains what can persist without conceding that content differences drive binding differences. The design takes three things from this debate: legitimacy for both sides of its test; the parsimony rationale for which side gets the null; and D7's feature list of stratified eschatology, covenant density, embodiment, and persistent bonds as its case-driven analogue of Norenzayan's move, naming structural content features in advance and asking whether they covary with binding. Even the study's category term is positioned within a live dispute: calling living traditions "mythologies" is defined structurally (narrative + ritual prescription + sacred/profane boundary), with Asad's (1993) counter to Geertz (1973), attend to authorization, discipline, and access, concretely absorbed as D1 and D4 rather than adjudicated, and Taylor (2007) framing the low-transcendence comparators as immanent-frame options whose ABSENT social codes embarrass nothing.

The two motivating commitments from §1 sit at the two poles of the crossing. The founding move of the bracketing argument, evaluate traditions by their premortem life-hypotheses without adjudicating the untestable, is methodologically Durkheimian (function without metaphysical adjudication) while its evaluative demand for lucidity under uncertainty is Camusian. The content prediction supplies what the social axis tests. The design is the disciplined collision of the two.

***Plain-language guide: how the two frameworks cross***

Put simply, Durkheim asks, "What helps a group stay together?" Camus asks, "What happens when people want a final meaning but cannot prove one?" The study examines these questions on separate tracks. It did not test whether stronger communities require less open uncertainty, whether they can support more of it, or whether the two are unrelated. That joint question is RQ3, and it remains untested.

**2.4 The registered research questions**

The design registers three questions the hypotheses serve. RQ1, the Durkheimian question: which structural features of a mythology's life-hypothesis, if any, covary with integration and regulation outcomes when intensity of collective practice is held comparable? RQ2, the Camusian question: what absurd-resolution strategy does each mythology employ, and at what lucidity cost, as evidenced in its canonical texts and prescribed practices? RQ3, the joint question: how do binding power and lucidity relate across cases, as trade-off, independence, or synergy? (Whether the gathering points beyond itself or *is* the thing gathered around, the question that closed the motivating content prediction, is a special case of RQ3.)

In plain language, the three questions are:

- **RQ1:** When traditions have similar levels of shared practice, do differences in their teachings help explain how strongly their communities hold together and guide members?
- **RQ2:** How does each tradition respond when people want a final meaning but the world gives no answer they can prove? Does the response keep that uncertainty open, close it with hope, or remain unclear in the available record? What, if anything, must people be asked not to examine closely for that answer to hold?

- **RQ3:** Does a group become stronger only when it gives members firm answers, can strong belonging and open uncertainty exist together, or can each change without the other? The present study did not test this question.

This record answers RQ1 with a disciplined null (§4.1), RQ2 with a per-tradition strategy classification and an argued judgment (§4.2), and leaves RQ3, including its registered trade-off, independence, and synergy alternatives, open work (§5.1, §6).

### 2.5 The registered hypotheses

Displayed in full, each with its ancestry:

A brief guide to the labels helps prevent a common mistake. A **null hypothesis** is the starting claim the study tries to find evidence against. An **alternative hypothesis** states another possibility. **Not rejected** means the study did not find enough evidence against the null; it does not prove the null is true. **Rejected (argued)** means a text-based, contestable judgment met the registered rule; it is not a statistical test. **Untested** means the study lacked the needed evidence or analysis.

- **$H_0$ (composite, as registered):** "the content of a faith tradition's mythology is causally and existentially inert." (The umbrella null over the two subnulls below, which alone carry registered rejection conditions; its joint survival is itself a registered finding, the construct §4.2 reports as not triggered.)

  *Plain-language gist:* The story a faith teaches may make no difference to either social binding or the way it handles unanswered meaning, within the parts of those ideas this study tests.

- **$H_0$-social (the Durkheimian null):** all binding work is done by collective practice itself, by gathering, ritual, and a maintained sacred/profane boundary; any mythology enacted at

comparable intensity binds and regulates comparably; the story is interchangeable, and only the telling-together matters. (Strict Durkheim; the practice pole of the debate above.)

*Plain-language gist:* The starting claim says shared practice does all the work of holding a group together and guiding its members. If practice levels are similar, the exact story should not explain a difference.

- **$H_1$-social:** specific content structures, namely stratified eschatology, covenant density, embodiment doctrine, and persistent kinship bonds, predict differential binding. (The motivating content prediction in its Durkheimian reading; the content pole, in Norenzayan's covariational form.)

  *Plain-language gist:* Specific teachings about the afterlife, promises, bodies, and family bonds may help some traditions bind differently from others.

- **$H_2$-social (Camus's rival):** durable solidarity can arise from shared revolt without sacred apparatus (Camus 1991b).

  *Plain-language gist:* A lasting community may form through a shared refusal to give up or claim a final answer, even without sacred beliefs or rituals.

- **$H_0$-existential (strict Camus, as registered):** "Every religious mythology resolves the absurd by the same move — philosophical suicide. Differences between traditions are stylistic, not structural. No mythology sustains lucidity; all trade it for hope."

  *Plain-language gist:* The null has two readings. On the narrower reading, every religious mythology may close the clash through hope. On the broader reading, no mythology of any kind may keep the uncertainty fully open. The study reports both because they lead to different conclusions.

- **$H_1$-existential:** at least one tradition's canonical texts and prescribed practice sustain the confrontation without asserting that the silence answers, a resolution differing *in kind* from the leap. The design registers in advance that $H_0$-existential is decidable only by *argued interpretive judgment*, not mechanical test, and that the codebook's descriptive existential categories presuppose the kind-distinction and therefore cannot serve as warrant (the registered C2 circularity).

  *Plain-language gist:* At least one tradition may face the unanswered question without claiming that it has been solved. Evidence for this difference must come from more than the labels used to sort the texts.

- **$H_2$-existential:** stable practice-with-professed-nonassent, members running the ritual engine at full binding power while publicly professing nonassent to the death-hypothesis (public profession being the codeable datum; Chaves 2010).

  *Plain-language gist:* A community may keep its rituals and social ties working strongly over time even when members publicly say they do not accept its claims about death. This study could not infer anyone's private belief.

Of these, this record returns dispositions on the two nulls (§4): $H_0$-social NOT REJECTED, and $H_0$-existential rejected only on its universal reading while not rejected on its religious-mythology reading, per the accepted dual reading; $H_1$-social found no support at the study's scale; $H_1$-existential was carried only by the Camus comparator under that dual reading; $H_2$-social and $H_2$-existential remain registered and untested.

### 2.6 The rejection conditions, fixed in advance

$H_0$-social is rejected only if, across cases *matched on practice intensity*, structural content features systematically covary with integration/regulation outcomes. $H_0$-existential is rejected

only by an argued, text-anchored, contestable judgment meeting the design's "kind, not decoration" standard. Both conditions were registered before any coding existed; the QCA's outcome rules were additionally frozen in a no-outcomes preregistration before any outcome value was seen (§3.5).

## 3. Method

### 3.1 Role separation and provenance controls

The reason for role separation is straightforward: the author is not outside the material he selected and framed. Waddell is a Latter-day Saint and reports greater familiarity with LDS doctrine and official sources than with the other cases. The accepted design left operator positionality to the operator to declare or withhold; this paper exceeds that registered minimum by making the disclosure part of its reflexivity treatment. The disclosure does not show that any choice was biased, but it identifies where scrutiny belongs: the research questions, framework, dimensions, case set, corpora, and outcome mappings were author-governed choices that cross-family coding could not make independent.

In plain terms, the provenance apparatus lets a reader ask four questions: can the process be retraced, can claims be tied to a record rather than a preference, did independent coding families reproduce the same judgments on the supplied material, and can another researcher see what would have to travel to a new study? More precisely, the decision ledger, sealed keys, and hash-locked submissions form a **dependability** audit trail. Per-cell provenance, locator confirmation before extraction, and preregistration before outcomes support **confirmability**. Blind double-coding across model families, with accepted agreement of 70/108 exact and 91/108 substantive, contributes to **credibility** conditional on the supplied corpora. The specified case

set, codebook, and partition tables support assessment of **transferability** by making scope and operational choices explicit; they do not establish generalizability or cross-case equivalence.

The study's methodological contribution is a provenance design for AI-assisted qualitative and configurational research. Waddell is the sole author and retained sole human decision authority. Fable 5 and Opus 5 assisted with design, instrument, manuscript, and analysis drafting; GLM and Nemotron coded; DeepSeek reviewed and base-adjudicated; Kimi readjudicated and reviewed; and gpt-5.6 verified sources, drafted the attributed existential judgment, and assisted with manuscript finalization. These roles conferred neither authorship nor decision authority. Independence was task-specific, not global.

The design used five controls: one versioned project record isolated from cross-project and provider-level memory; an immutable decision ledger that separates authorization, execution, evidence, and acceptance; scoped hash attestation of protected and analytic artifacts; role separation with every exposure disclosed; and preservation of attributed disagreement and repair history. The Anthropic drafting family was excluded from coding and adjudication. The coders were blind to the calibration key and to each other's submissions; GLM's prior design exposure remained disclosed. The original calibration key was sealed before any coder was designated, and an operator-authorized correction was made only after both calibration submissions were locked and was then resealed. Same-family-assisted analysis and every later role substitution remained visible in the record.

The source-equivalence control was limited. A canon review performed outside the drafting model family preceded coding and identified gaps that were repaired or explicitly routed, but the P5 coders then applied the same pre-enumerated per-case corpora rather than constructing independent corpora. Their agreement therefore bears on coding reproducibility

conditional on those inputs. It does not establish equal corpus completeness, representativeness, or access across cases.

These controls establish sequence, byte identity where pinned, and failure visibility. They do not validate a substantive claim, make model families globally independent, control the author's framing choices, or show that this workflow outperforms conventional research production. Principal records, access status, and attestation scope appear in "Data, Materials, and Audit-Record Availability"; the complete governed inventory is maintained in the workspace record described there.

### 3.2 Instrument

The approved second version of the codebook operationalizes the design: the Durkheimian-axis dimensions D1–D7, with D1–D6 covering the social and practice structure and D7 carrying the partition-neutral structural content features that become QCA conditions, and the existential categories C1–C5 (C1 absurd acknowledgment; C2 resolution type, descriptive only; C3 lucidity cost; C4 status of hope; C5 solidarity source), coded from public and professed sources only under the codebook's rule P-5, "no private states": codeable data are public representations, meaning texts, prescriptions, professed statements, and documented behavior, following Chaves 2010 on the religious congruence fallacy. Accepted per-case claims partition tables separate testable life-hypotheses from bracketed death-hypotheses for every case, with every row cited.

### 3.3 Case set

Nine cases, fixed by an operator-accepted census rule over the Pew global frame at ≥ 0.3 billion adherents meeting a canon criterion (Hackett et al. 2025): five high-content/high-practice

traditions, the Latter-day Saints, Nicene Christianity, Sunni Islam, Gītā-anchored Hinduism, and Theravāda Buddhism, and four registered comparators completing the factorial contrast: materialist-panpsychism ($H_4$; Strawson 2006; Goff 2019), the simulation hypothesis ($H_5$; Bostrom 2003), secular civic ritual, and Camusian absurdism (Camus 1991a), the last deliberately included as "the Camus-null embodied" under a disclosed category-stretching concession carried as a named threat (§5.3).

### 3.4 Calibration, blind coding, and adjudication

The role-selection rule was constraint first, preference second. In plain terms, a family was first screened for authorship, key access, prior review, and the independence required by the task. Only when more than one eligible family remained did behavior observed in this workspace break a tie. Those observations were task-specific; the record makes no claim about any model's general capabilities.

The early instrument review illustrates the ordering. DeepSeek held no O2 authorship and had found six material gaps on the first O1 pass, so its local review record supported a role for which its independence already made it eligible. The blind-coder pool was then quarantined until after the scenario key was sealed. Its task-local rationale paired Kimi's observed faithful execution of bounded instructions with GLM's observed careful fresh reading; the quarantine and later exposure rules remained controlling. For real-case coding, GLM and Nemotron were eligible because neither family had authored, edited, or reviewed Codebook v2, and both had passed sealed-key calibration. The calibration supplied local differentiation: GLM matched the adjudicated answer more closely on all thirteen contested rows, while Nemotron was clean on uncontested rows and identified the hidden S5 C3 trap. These are recorded performances on this instrument, not general temperament or vendor claims.

The nine cases were then coded independently by GLM and Nemotron, blind to the calibration key and to each other's submissions, with GLM's prior review of the accepted design disclosed, under a registered source-access protocol; submissions were hash-locked on filing. DeepSeek base-adjudicated all 108 combined D- and C-dimension cells (69/108 ≈ 64% exact, 91/108 ≈ 84% substantive); a repair of fabrication-affected rows in one submission was readjudicated by Kimi, yielding the accepted rates of 70/108 ≈ 65% exact and 91/108 ≈ 84% substantive, with a counting-unit ambiguity later resolved by operator ruling.

Later assignments kept tradeoffs visible rather than translating them into rankings. Opus supplied continuity for the QCA while the record named its same-family authorship as the weakest independence option and required a no-outcomes registration plus non-Anthropic review. For the existential judgment, the operator chose gpt-5.6 after a comparison that separately named Kimi's clean hands and literature exposure, Opus's author-led continuity, DeepSeek's preserved reviewer lane, and gpt-5.6's prior verification exposure. Kimi's later paper-review role was tied partly to checking exact carriage of the qualification it had authored. For R8, Kimi rather than DeepSeek was reserved because this selection account repeats the ledger's strongest favorable project-specific assessment of DeepSeek. That choice manages the narrower conflict; it does not rank either family. Each rationale is bounded to the recorded task and exposure.

The D7 anchors for three cases then went through an independent source-verification pass (gpt-5.6, VF-01..VF-22), discharged with 22 adjudicated findings, 21 UPHELD and one MODIFIED, moving no D7 value, with one anchor accepted expressly unverified (P19; Fuller 2004, ch. 9) and eight negative/summary anchors carried into every later table as limitations under the codebook's rule P-4, "absence is a value": a corpus-wide ABSENT or summary code

is legitimate, but a fetched passage cannot prove a whole-corpus negative. The record was accepted as *the coded data, not a finding*.

### 3.5 Preregistration of the QCA

Before any outcome source was opened, the designated analyst (Opus 5, same-family-assisted and disclosed) filed a no-outcomes registration fixing: four dichotomous conditions from D7 (STRAT, stratified eschatology; COVD, covenant density; EMB, embodiment doctrine; PBOND, persistent postmortem kinship bonds; the fifth D7 feature excluded in advance on circularity grounds); the binding outcome, tradition-level retention, the share of those raised in a tradition who still identify with it, from open Pew Religious Landscape Study 2023–24 outputs (Smith et al. 2025); a practice-intensity matching rule (weekly attendance, median ± 20 points, flag-and-report); comparators as structurally unavailable design cells (N = 5 computed); BIND = retention ≥ 0.50 with a binding null-variance clause; five sensitivities; and a minimal analytic ceiling, with QCA mechanics per Ragin (1987, 2008), consistency and coverage benchmark practice per Schneider and Wagemann (2012), and the configurational-comparative procedures in Rihoux and Ragin (2009). A non-Anthropic review (DeepSeek: no material gap), six advisory repairs verified by the reviewer, and operator rulings on every open item closed the registration before the truth table was authorized.

### 3.6 Analysis and review chain

The truth table was built under the frozen registration, then dual-reviewed by independent passes, DeepSeek recomputing every statistic in-workspace (NO MATERIAL GAP; the $H_0$ reading HOLDS) and gpt-5.6 refetching every published value from the sources (ALL VALUES VERIFIED), with zero conflicts between the passes, five advisory repairs verified in place, and

operator acceptance. The existential judgment (§4.2) was drafted by gpt-5.6 from a 36-cell adjudicated reconstruction and reviewed by Kimi (36/36 cells confirmed; qualified concurrence), and accepted with its qualification.

## 4. Results

Figure 1 is a map of the registered design and the two accepted analytic dispositions. It keeps the evidence streams separate and places their inferential ceilings before the detailed results. The connectors show document hierarchy only; none represents causation or an association between the axes.

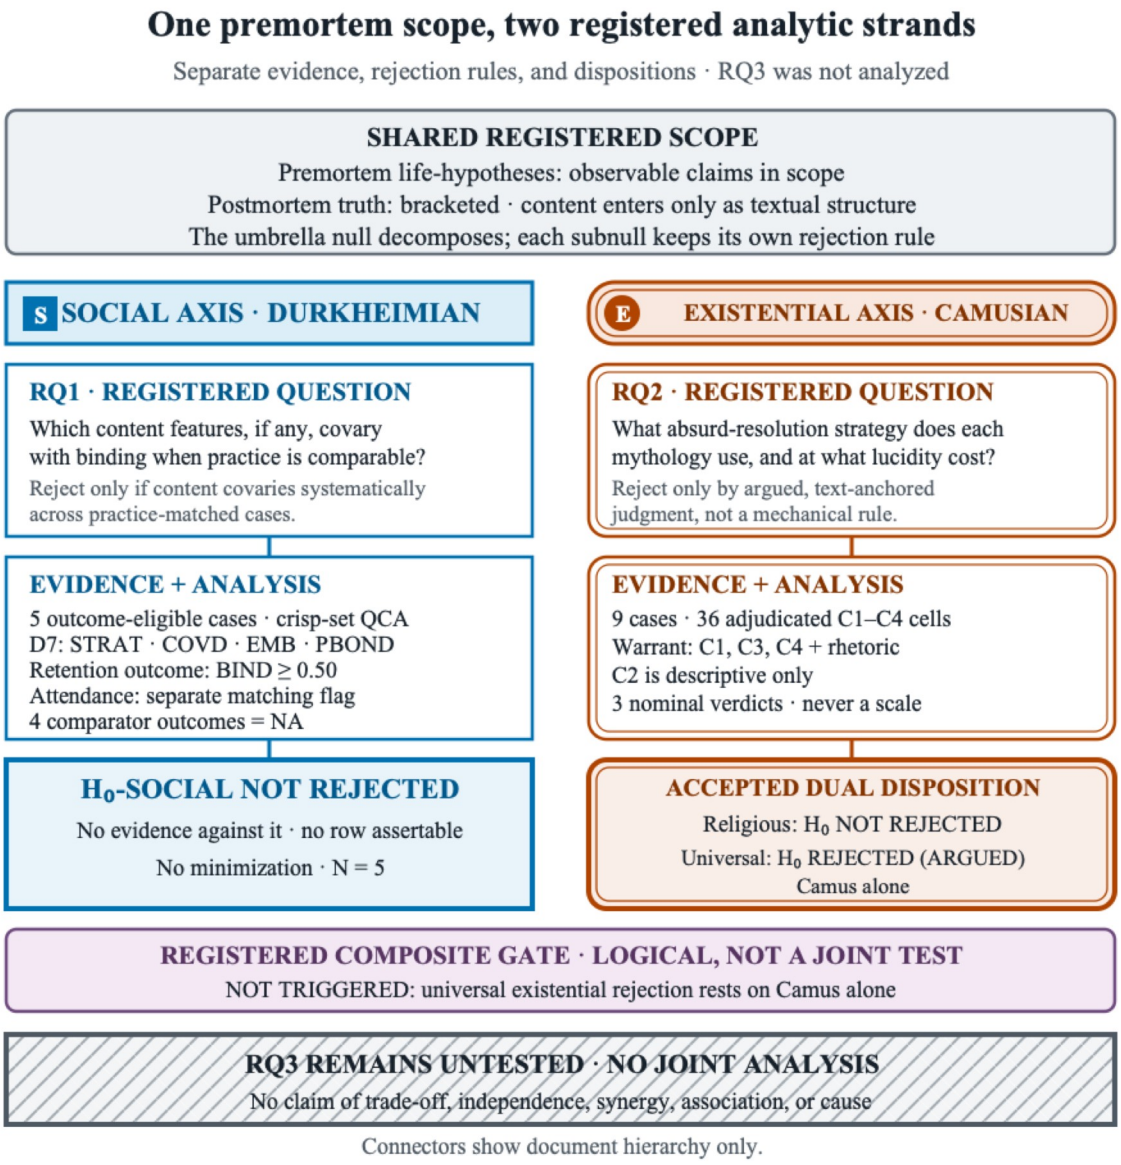


Figure 1. Conceptual diagram of the registered two-axis design. A shared premortem scope leads to separate social and existential lanes. The social lane ends in $H_0$-social not rejected after a five-case QCA. The existential lane ends in religious-mythology nonrejection and universal rejection carried solely by Camus after a nine-case argued judgment. The lanes do not merge. A bottom barrier states that the registered composite condition did not trigger and RQ3 remains untested because no joint analysis was performed.

*Source:* the registered research design, the frozen QCA preregistration, the accepted truth-table analysis, and the accepted existential judgment and its cross-family review (see "Data, Materials, and Audit-Record Availability"). Visualization generated by the study's deterministic sixth-revision figure script.

*Note:* The boxes reproduce registered components and accepted dispositions; they add no estimate or analysis. The lanes use different evidence, units, and decision forms and must not be read as a joint comparison. NOT REJECTED means no evidence against the null at this study's scale, not proof that the null is true or that effects are absent. Existential verdicts are nominal; UNDECIDABLE-ON-RECORD is not an intermediate score. The composite gate is logical, not a test of the relationship between binding and lucidity. RQ3 remains untested.

### 4.1 The social axis: no evidence against $H_0$-social

Under the registered decision rule, $H_0$-social was not rejected: the analysis produced no evidence against it. The values and configurations below explain that disposition.

The published retention values (Smith et al. 2025, 109; verified against the source by an independent pass) and the frozen ≥ 0.50 cut gave: LDS 54% (BIND = 1), Nicene via the flagged Catholic proxy 57% (1), Sunni via the Muslim category 77% (1), Hinduism via the Hindu category 82% (1), Theravāda via the Buddhist category 45% (0). Weekly-attendance values (Smith et al. 2025, 166) produced 69, 29, 39, 18, and 6 (median 29; band [9, 49]), flagging LDS (high) and Theravāda (low) as matching failures, retained with flags per the registered rule.

Figure 2 places the outcome and matching variables on separate spectrum displays. Threshold zones, direct labels, and distinct marker shapes make the registered cuts and flags visible without combining the measures. The alignment is descriptive; neither panel identifies what caused binding.

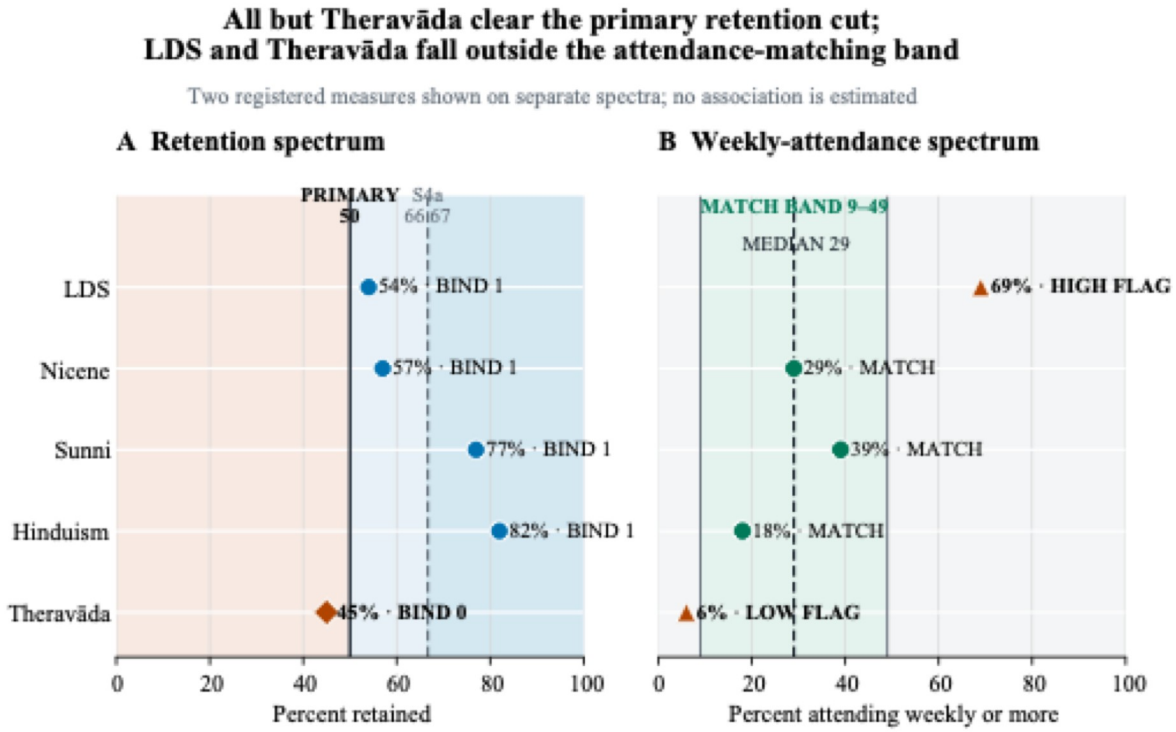


Figure 2. Separate retention and weekly-attendance spectra for the five outcome-eligible cases. All but Theravāda clear the primary retention cut; LDS and Theravāda fall outside the attendance-matching band.

*Source:* the accepted truth-table analysis; published values from Smith et al. (2025, 109, 166). Derived from the study's frozen master visual under sixth-revision numbering, accessibility additions only, no plotted content changed.

*Note:* N = 5, US only. LDS uses the exact source category; Nicene Christianity, Sunni Islam, Gītā-anchored Hinduism, and Theravāda Buddhism use the Catholic, Muslim, Hindu, and Buddhist proxy or parent categories respectively. The 50 percent line is the primary BIND cut; 66.67 percent is sensitivity S4a. The attendance band is the registered 9–49 percent flag-and-report rule around the five-case median of 29. Both flagged cases remain in the primary analysis.

The truth table produced four observed configurations. The three distinct content vectors, LDS (STRAT·COVD·EMB·PBOND), Nicene (¬STRAT·COVD·EMB·¬PBOND), and Sunni (STRAT·¬COVD·EMB·¬PBOND), are each single-case rows with the same outcome (BIND = 1), each at consistency 1.00 and raw coverage 0.25: *content configuration does not sort the outcome among them*. The shared all-zero vector (Hinduism, Theravāda) is a contradictory row (outcomes 1 and 0, consistency 0.50), at exactly the site the registration predicted a contradiction would appear if one appeared at all. No minimization was possible; the complex solution is three single-case terms (solution consistency 1.00, solution coverage 0.75). No row is assertable as a

finding, on two grounds fixed in the registration and stated separately: rows A–C are n = 1 rows whose consistency of 1.00 the registration reports in advance as "consistent but uninformative," never support for a hypothesis, and row D sits below the registered 0.8 consistency bar and is reported as a contradiction, never resolved.

Figure 3 makes the shared content vector and split outcome in row D visible; Table 1 then supplies the exact audit values and the registered with-and-without checks for the two practice-matching flags.

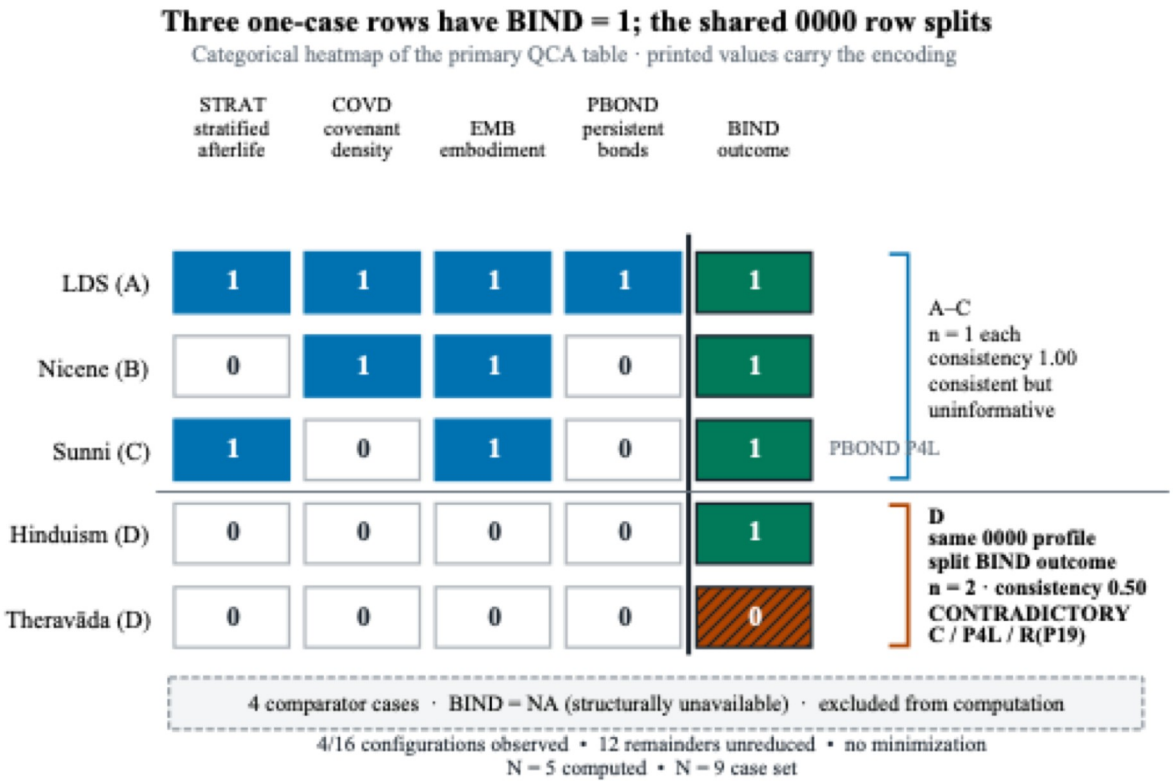


Figure 3. Categorical heatmap of the primary QCA table. Three one-case rows have BIND = 1; the shared 0000 row splits and is contradictory.

*Source:* the accepted truth-table analysis. Derived from the study's frozen master visual under sixth-revision numbering, accessibility additions, one headline clarification, and punctuation conformance, no plotted content changed.

*Note:* Rows A–C are n = 1 and consistent but uninformative; row D is contradictory. The chart uses the primary BIND calibration. It excludes the four structurally unavailable comparator outcomes, leaves twelve remainders unreduced, and does not demonstrate a configurational or causal effect.

**Table 1. Primary QCA truth table and registered with-and-without checks**

**Panel A. Observed configurations**

| Row | Vector | Case(s) and BIND | n | Cons. | Registered reading |
|---|---|---|---|---|---|
| A | 1111 | LDS (1) | 1 | 1.00 | Consistent but uninformative |
| B | 0110 | Nicene Christianity (1) | 1 | 1.00 | Consistent but uninformative |
| C | 1010 | Sunni Islam (1) | 1 | 1.00 | Consistent but uninformative |
| D | 0000 | Hinduism (1); Theravāda (0) | 2 | 0.50 | **Contradictory** |

**Panel B. With-and-without the two matching-flagged cases**

| Variant | N | BIND = 1 | Contradiction | Coverage | Registered reading |
|---|---|---|---|---|---|
| Primary (both retained) | 5 | 4 | Yes (row D) | 0.75 | Three single-case terms |
| Excluding LDS | 4 | 3 | Yes (row D) | 0.67 | Two terms; PBOND has no variance |
| Excluding Theravāda | 4 | 4 | No | 1.00 | **Null variance; nondiscriminating** |
| Excluding both | 3 | 3 | No | 1.00 | **Null variance; nondiscriminating** |

*Source:* the accepted truth-table analysis.

*Note:* In Panel A, vector order is STRAT, COVD, EMB, PBOND, and every row has raw coverage 0.25. The LDS attendance value is match-flagged; Nicene and Sunni use Catholic and Muslim proxy or parent categories; row C carries PBOND P4L; row D carries dependencies C, P4L, and R (P19), grade codes defined under "Scope and carrying conditions" below, and Theravāda is match-flagged. Four of sixteen configurations are observed; twelve logical remainders remain unreduced. The four comparator outcomes are NA (structurally unavailable) and excluded from every computation. No row is assertable as a finding.

All three registered sustained-null signatures are present, and the report's sharpest sentence survives an independent arithmetic recomputation and an independent source-value verification: **excluding Theravāda, one of the two matching-flagged cases, measured on the weakest proxy of the five, removes all outcome variance from the study.** The sole BIND = 0 case is also the case the registered matching rule says is not practice-matched. Sensitivity S4a (a source-mean threshold) flips the two marginal cases (LDS, Nicene) and is reported beside the primary; two registered sensitivities (S4b, S5) proved unexecutable because the sources lack the registered constructs, and were reported rather than repaired.

***Scope and carrying conditions***

This disposition carries the following scope and conditions as part of the result, not as supplementary caveats, exactly as the accepted analysis records them. The accepted record states at the registered altitude: **$H_0$-social is NOT REJECTED; this analysis produces no evidence against it.**

**Analytic and measurement scope.** N = 5 for every computed statistic (N = 9 for the case set). The four comparator outcomes are printed as NA (structurally unavailable) design cells and excluded from every computation. Four of five outcomes use proxy or parent categories; only LDS matches its source category exactly. LDS and Theravāda remain in the primary analysis with practice-matching flags.

**Source-record ceilings.** Every condition cell carries the verification grade of its anchor, registered before analysis and never dropped for readability: V marks a locator verified on a fetched text; C marks a locator corrected under an accepted adjudication with no value moved; P4L marks a corpus-wide ABSENT or summary code accepted as a stated limitation under the codebook's P-4 rule; and R marks an accepted unverified residual. ([INF], in §4.2, marks a code carried as coder inference rather than quoted text: admissible under the codebook's P-2 evidence standard, though no cross-case claim may rest on [INF] codes alone.) The grade dependencies travel row by row. Row C's PBOND cell is P4L. Row D carries C, P4L, and the sole accepted residual R (P19, Nemotron's Hindu persistent-bond anchor, Fuller 2004, ch. 9). That residual bears on the anchor rather than on any contested value: both coders coded the feature absent, and GLM's independent P4L row (P14) corroborates it. The eight negative or summary P-4 anchors (P04, P09, P12, P14, P22, P24, P27, P29) are treated neither as verified nor as missing data because a fetched passage cannot prove a whole-corpus negative. The three locator corrections (VF-01, VF-03, VF-04) are recorded only in the decision ledger and move no D7 value. Tier-2

sources (Besant 1922; Sujato 2018a, 2018b, 2018c; Ṭhānissaro Bhikkhu n.d.) establish structure only, never held-edition wording.

**Execution and scope limits.** The US-only scope could not be tested globally because sensitivity S5 was unexecutable; the regulation-composite sensitivity S4b was likewise unexecutable on this wave. The analyst disclosed two execution defects: an incomplete HTML-to-Markdown rendering, a locator-reliability failure self-caught and corrected in a filed revision, and incidental value exposure before locator filing, with all thresholds frozen beforehand. The analyst also disclosed instrument exposure through same-family-assisted work. The Hindu D1–D6 record is secondary-sourced (Fuller 2004) and has the widest text–practice gap of the nine. The Sunni matching note retains its recorded repair history; it now reads "salat's prescribed daily frequency."

**Design threats.** The design's complete seven-item threat set remains: functionalist circularity, guarded by defining mythology structurally and measuring binding strictly as an outcome (§2.1) and by excluding the premortem-consequence feature from the registered conditions in advance; survivorship bias, because extant traditions are the winners of a selection process whose failures are unobserved; confounding, because content never varies independently of history, ethnicity, economics, and persecution, so that EMB, the highest-coverage condition, coincides with an Abrahamic-plus-LDS versus Indic split; normatively loaded Camusian codes; category stretching at the comparators; coder bias at corpus scale, because all current model families share substantial corpus overlap; and the design's own recursion, because the design is itself a meaning-making artifact produced under the conditions it studies. The registered three-vector limit also remains. No decision rule fired; no significance statistic exists anywhere in the analysis; no row is assertable as a finding.

### 4.2 The existential axis: universal rejection, religious-case nonrejection

The accepted disposition is dual and governs every case classification that follows. Kimi's review confirmed conformance with no material gap and entered a **qualified concurrence** that the operator's acceptance made permanent:

> On the registered null's universal sentence ("no mythology sustains lucidity"), **$H_0$-existential is REJECTED (ARGUED)**—carried solely by the Camus comparator. On the null's religious-mythology head-noun reading, **$H_0$-existential is NOT REJECTED on this record**—every decidable religious case is LEAP, and the two remaining religious candidates are undecidable pending an operator-authorizable text-opening extension.

The evidence for that disposition comes from the adjudicated C-record: all 36 cells were reconstructed with per-cell provenance and independently confirmed 36/36 by review. The argued judgment (gpt-5.6) applied a protective discipline, under which a case failing the necessary C1 gate is *undecidable*, not a leap, absent affirmative evidence:

- **LEAP**: LDS, Nicene Christianity, Sunni Islam. In each, the corpus stages or acknowledges no unanswered silence it does not answer: mandated hope, carried at [INF] grade for the LDS and Sunni cells per the adjudicated record, closes the confrontation (LDS covenantal assurance; "hope … does not disappoint," Rom 5:5, and "the blessed hope," Titus 2:13, in the Nicene record, whose C4 cell carries an unresolved anchor-grade and column-scope split between the coders alongside its edition looseness; the Garden and divine pleasure in the Sunni record, the hope classification an adjudicated inference on verified locators). Notably, the strongest religious candidate (Nicene Christianity, whose corpus *dramatizes* the confrontation in the cry of forsakenness) is a leap precisely because the same corpus answers it.
- **UNDECIDABLE-ON-RECORD**: Gītā-anchored Hinduism and Theravāda Buddhism (each missing decisive C3/C4 cells, their C1/C2 cells carried at [INF] grade with Tier-2

structure-only verification; Theravāda is the strongest religious third-move candidate, MN 72's refusal being stable across two translations, but its only proposed hope anchor failed source verification and nothing replaced it), and the three non-Camusian comparators (no coded confrontation to judge).

- **SUSTAINS**: Camusian absurdism. C1 is named and C4 is coded dissolved because the recorded prescription refuses suicide and hope; both are carried at [INF] grade on offline-edition anchors, with C3 empty after an explicit antideflection look-for, and the prescribed conduct itself read as the practice bridge, a reading the judgment marks expressly contestable and load-bearing.

Figure 4 displays the accepted verdicts as separate nominal cards by case class, then colors the two fixed dual-reading phrases without sizing or positioning them by frequency. UNDECIDABLE-ON-RECORD is not an intermediate score or a weaker form of LEAP.

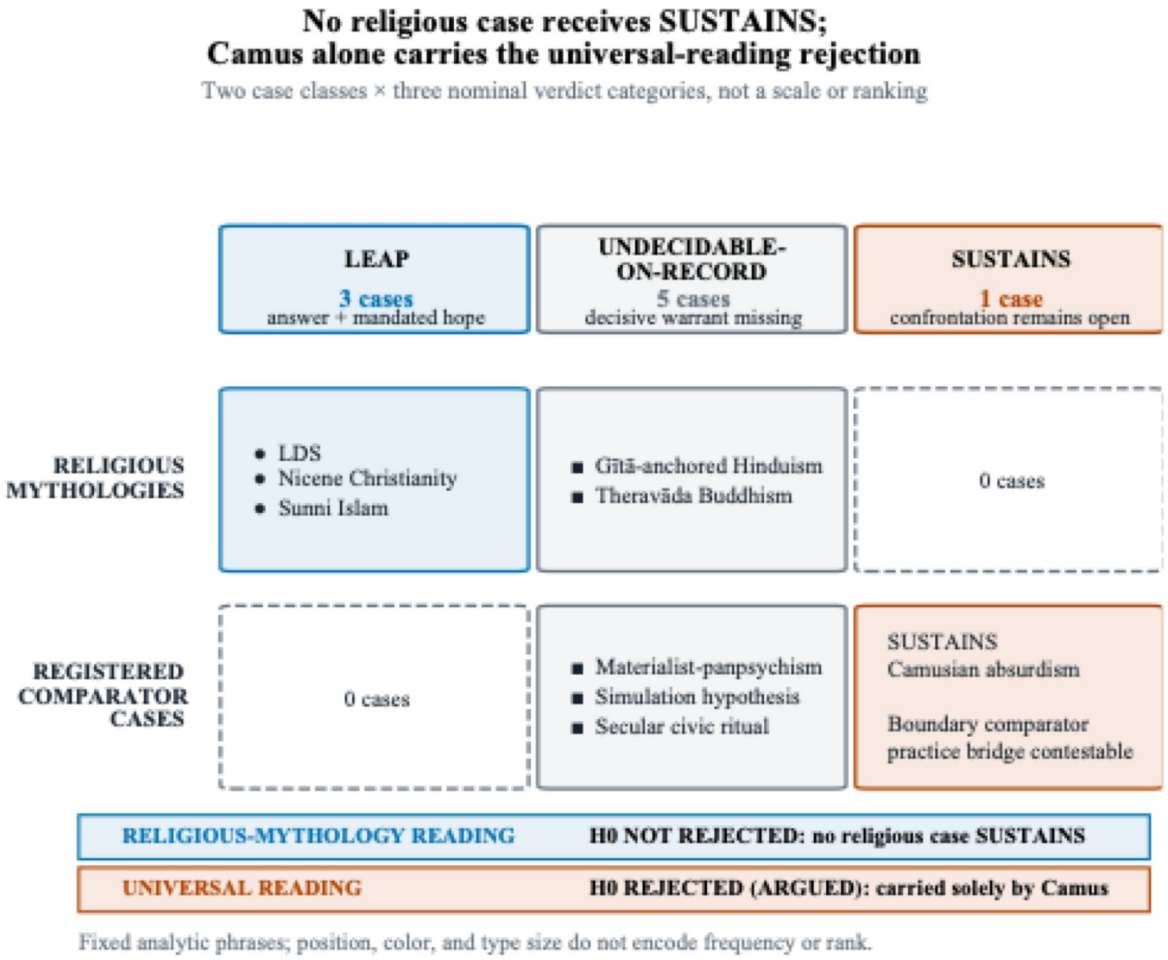


Figure 4. Nominal verdict-card matrix by case class: three LEAP, five UNDECIDABLE-ON-RECORD, and one SUSTAINS. No religious case receives SUSTAINS; Camus alone carries the universal-reading rejection.

*Source:* the accepted existential judgment and its cross-family review, accepted with qualification. Derived from the study's frozen master visual under sixth-revision numbering, accessibility additions and punctuation conformance, no plotted content changed.

*Note:* Verdicts are nominal, not ordinal. No religious case received SUSTAINS; Camusian absurdism alone carries the universal-reading rejection. On the religious-mythology reading, $H_0$-existential is not rejected, and the Camus practice-bridge reading remains contestable.

The reviewer's formulation is preserved verbatim in the record: the rejection is "tautology-adjacent" as evidence about religion, because the null is Camus's own categorical verdict, the sustaining case is that verdict coded as a case, and so the rejection is true substantially by construction, though it is *earned* at one point (the empty C3 row is a coded, two-coder, adjudicated datum, not a tautology). Because $H_0$-existential was rejected on the universal reading, the design's registered both-nulls finding ("functionally interchangeable binding engines

whose specific content is existential anesthesia") was not triggered; a single comparator's verdict blocked it.

### *Scope and carrying conditions*

This disposition carries the following scope and conditions as part of the result, not as supplementary caveats; each restates a limit already recorded in the accepted judgment, its review, or the acceptance itself.

**Warrant and instrument scope.** The judgment is an argued, text-anchored, contestable interpretation under the design's analytic ceiling; "kind, not decoration" has no mechanical rule. The registered C2 circularity stands: the codebook's descriptive resolution-type categories presuppose the kind-distinction the judgment argues, were barred as warrant, and were policed at review: every C2 value is description only.

**Source-record ceilings.** The verdict-bearing cells carry their coding-record grades (defined in §4.1). The LDS and Sunni mandated-hope cells are carried at [INF] grade, and the Nicene C4 cell carries an unresolved anchor-grade and column-scope split between the coders alongside its edition looseness. The Hindu and Theravāda C1/C2 cells are carried at [INF] grade with Tier-2 structure-only verification, and each case is missing decisive C3/C4 cells, which is why both are UNDECIDABLE-ON-RECORD rather than classified. Camus's C1 and C4 cells are carried at [INF] grade on offline-edition anchors, and its empty C3 row follows an explicit antideflection look-for, a coded, two-coder, adjudicated datum.

**Interpretive dependencies.** The universal-reading rejection is carried solely by the Camus comparator; the practice-bridge reading of its prescribed conduct is disclosed, load-bearing, and expressly contestable; and the rejection is "tautology-adjacent" as evidence about religion. Verdicts are nominal, not ordinal; UNDECIDABLE-ON-RECORD is not an

intermediate score. The two undecidable religious cases remain decidable under the operator-authorizable text-opening extension (§6), without reopening the accepted judgment.

## 5. Discussion

### 5.1 A cross-axis synthesis: attributed, not a third analysis

What follows is the paper's synthesis, drafted with Fable 5 assistance and presented under the author's responsibility. The record contains the QCA result and the existential judgment. Nothing in the record performs a joint analysis, so RQ3 remains untested. Read separately at their registered altitudes, the two results do not conflict; that observation is synthesis, not an RQ3 test.

On the social axis, no evidence against $H_0$-social emerged: the three distinct content configurations each had the calibrated outcome BIND = 1, the shared configuration split, and the only observed outcome variation aligned with the practice-matching flag, a nonrejection with a suggestive alignment, not an affirmative finding that practice binds. On the existential axis, all three decidable religious cases answer the silence rather than sustaining it, with mandated hope as the coded mechanism; two religious cases remain undecidable, and the only SUSTAINS belongs to the posture that refuses religious form altogether. The analysis did not reject the Durkheimian null; among decidable religious cases, the argued judgment likewise did not reject the religious-mythology reading of the Camusian null, which fell only on its universal wording through its registered comparator. The two motivating propositions meet in the middle: the demand that traditions be judged by premortem life-hypotheses was executed as a research design, while the content-binding prediction ($H_1$-social) found no support at this N, though the study's power to detect it was, by registered design, minimal.

### 5.2 What the Camus boundary result is worth

The narrow rejection is the record's most philosophically instructive artifact. Had the case set held only religious traditions, $H_0$-existential would stand unrejected and the both-nulls finding would have fired. Including the null's own embodiment as a case turned the design into an instrument that can distinguish "no mythology sustains lucidity" (false on this record: one does) from "no *religious* mythology sustains lucidity" (unrefuted on this record). That distinction, invisible until a comparator carried a verdict, is offered here as a design warning and a hypothesis for future comparator-bearing studies, not as a validated general prescription: what this record demonstrates is that a comparator can shift a null's quantifier domain and manufacture a formally valid but low-information rejection, and that the danger is survivable only if the record refuses to resolve the readings into one, as this one does, with the analyst's verdict, the reviewer's qualification, and the operator's dual-reading acceptance all standing attributed.

### 5.3 Threats to validity

Author positionality and measurement asymmetry come first. As §3.1 explains, Waddell's LDS membership and greater familiarity with LDS doctrine and sources may have made that case easier to specify and audit while making omissions elsewhere harder to recognize. Cross-family coding controls coders; it does not control his choice of research questions, framework, dimensions, case set, or source architecture. The measurement record sharpens this concern: only LDS matched the outcome source's category exactly, full-text access varied by case, and the other four computed outcomes relied on proxy or parent categories. Corpus completeness and cross-case measurement equivalence remain unestablished. The

trustworthiness controls make this asymmetry inspectable. They are mitigation, not independence.

Scale, scope, and operationalization form one linked limit. With five computed cases, three content vectors, one contradiction, and US-only outcomes, the QCA can suggest but cannot estimate or identify causes. The weakest proxy carried all observed outcome variance, and the registered global sensitivity was unexecutable. Retention is a defensible observable trace of continued identification, not Durkheimian integration or regulation itself. Weekly attendance is observed participation, not prescribed frequency, and is cross-sectionally entangled with retention and unevenly comparable across traditions. Regulation remained untested, S4a changed two of five classifications, and selection, demography, migration, exogamy, category aggregation, reverse direction, civilizational history, ethnicity, economics, and persecution remain unexcluded.

Two interpretive limits remain. The C2 categories presuppose the kind distinction the existential judgment argues, so they were barred as warrant. “Kind, not decoration” has no mechanical rule; the judgment is an attributed, contestable argument, and its practice-bridge reading remains load-bearing and contestable.

AI production was concentrated in two families. Anthropic supplied most design, instrument, analysis, and initial manuscript-drafting assistance. OpenAI supplied the existential judgment and source-value verification, then concentrated manuscript finalization, figure production, builds, and this revision’s account of its own prior designations. Cross-family coding blind to the key and to the other coder, non-drafting-family adjudication, designated review, preservation of Kimi’s qualification, and hash attestation mitigate those concentrations. They do

not make the models neutral, remove shared training-corpus error, or turn author-side QA into independent review.

### 5.4 The apparatus as a contribution

Independently of its substantive findings, the record demonstrates that a contested humanistic question can be put through preregistration discipline usually reserved for quantitative fields, with the interpretive moments (the argued judgment, the qualified concurrence) not suppressed but *attributed and preserved*, using heterogeneous AI systems whose division of epistemic labor is enforced by protocol rather than trust: the initial calibration key sealed before coder designation, the later postlock repair and reseal disclosed, no-outcomes registration before source contact, adversarial review before every acceptance, and a verifier that re-proves the byte-identity of the pinned chain on demand.

The model-selection record adds a second methodological point. Heterogeneity alone is not independence. Roles were assigned by excluding families whose authorship, key access, or prior role conflicted with the task, and only then using performance observed on this workspace's own calibration and review record to choose among eligible options. Continuity, clean hands, reviewer preservation, and prior exposure were recorded as tradeoffs rather than collapsed into a claim that one model was generally stronger. That ordering is inspectable, including where access failures forced substitutions and where a reviewer was selected partly to verify carriage of its own prior qualification. R8's choice of Kimi over DeepSeek manages the passage's specific favorable account of DeepSeek rather than ranking the families.

The contribution claimed is bounded to what was demonstrated: traceability and failure visibility, not validation. Hashes establish byte identity, not truth; operator acceptance is governance, not empirical validation; model-family separation does not address error modes

shared through common training corpora; and no comparison condition was run, so nothing here shows the apparatus outperforms conventional production, only that the failure modes met along the way (a fabrication episode repaired under readjudication, an incomplete HTML-to-Markdown rendering self-caught and corrected, internal inconsistencies in summary prose surfaced by later reconstruction, access failures forcing role swaps, and this paper's own first-draft fidelity failures caught by its designated review) all left marks. The apparatus does not claim infallible components; it claims that component failures become visible.

## 6. Limitations and future work

Beyond §5.3, eight extensions remain open:

- A **corpus-equivalence audit** could ask, in plain terms, whether each case was represented by authoritative sources adequate to the defined case and coding questions. Paired insider and comparative scholars for each religious case, plus appropriate subject specialists for the comparators, would not be shown the study's outcome table or current codes. They could assess authoritative status, doctrinal and practice coverage, translation fit, and potentially code-changing omissions. Fresh coders could then report which cells and dispositions are sensitive to the expert-audited corpora.
- The **text-opening extension** specified in the judgment, opening the held Gītā and Pali materials, could decide the two undecidable religious cases and remains available without reopening the accepted judgment.
- **$H_2$-existential** (stable practice-with-professed-nonassent) remains untested and would require a member-practice corpus the current instrument does not collect.
- **$H_2$-social** (solidarity from shared revolt) was registered and remains untouched.

- A **wider case set**, including the deferred Judaism case or denominational disaggregation, could relieve the three-vector limit at the cost of the census rule.
- **Participation and retention over time.** In plain terms, which changes first: participation or continued identification? Longitudinal panel or cohort-linked data could examine whether higher participation precedes later retention, continued identification precedes later participation, or both track measured common causes. A future analysis would keep weekly attendance and retention separate. The present pilot supplies neither temporal ordering nor a causal finding.
- **Afterlife structure and participation.** In plain terms, is communal participation higher when a tradition portrays, in its own terms, a desirable outcome after death or release from rebirth as attainable? Does any relationship differ when attainment is tied to conduct, ritual, or authorized institutions? A future study would preregister separate codes for desirability, attainability, behavioral contingency, meaning whether conduct or ritual is made a condition, and institutional mediation, meaning whether access depends on authorized rites or offices. Institutional mediation would be defined independently of attendance frequency to avoid building the outcome into the predictor. The study would analyze prescribed participation separately from observed participation. The truth of each postmortem or soteriological claim would remain bracketed. This is a post-study hypothesis, not a finding from the present record: the pilot did not code those predictors, use attendance as an outcome, or analyze that relationship.
- A future study could source and preregister a norm-guidance measure before data opening, then execute a regulation composite prospectively.

## 7. Conclusion

Bracket what cannot be tested, and test what can. This pilot executed that founding partition with preregistered rules and a multimodel record auditable within the governed workspace; it does not show that mythological content is socially inert. On the social axis, $H_0$-social is not rejected. At N = 5, the analysis produced no evidence against the strict Durkheimian null, content-based binding found no support, and the only observed outcome variation aligned with the practice-matching flag. No row is assertable as a finding, and the study did not identify what caused binding.

The existential result must remain dual. $H_0$-existential is rejected (argued) on its universal reading solely because the registered Camusian comparator received SUSTAINS. It is not rejected on the religious-mythology reading: all three decidable religious cases are LEAP, two remain UNDECIDABLE-ON-RECORD, and no religious case received SUSTAINS. The comparator-carried rejection remains "tautology-adjacent" as evidence about religion, and the Camus practice-bridge interpretation remains contestable. The registered both-nulls finding therefore did not trigger, because the universal-reading rejection that fails its antecedent is carried by the Camus comparator alone. Because no joint analysis was performed, RQ3 remains untested; the two axes cannot support a claim of trade-off, independence, or synergy.

The methodological contribution is likewise bounded. Role-separated provenance made exposures, disagreements, repairs, and accepted qualifications traceable, while hash attestation made the pinned record auditable byte for byte. Those controls establish traceability and failure visibility, not empirical validation, freedom from shared model error, or superiority to conventional research production. At this small, US-only, proxy-dependent scale, the paper's contribution is a disciplined result and an inspectable route to it, not a general verdict on religion, binding, or lucidity.

## Data, Materials, and Audit-Record Availability

The public social-axis values are in Smith et al. (2025, 109, 166). All other study records are internal to *The Terminal Experiment* workspace, where stable filenames serve as locators; they are not a public repository or DOI, and they are audit record rather than citable sources. The principal records are the research design and null hypothesis (O1-Research_Design_and_Null_Hypothesis-20260816.md), the operative codebook (O2-Codebook_v2-20260817.md), the per-case claims partition tables (O2-Claims_Partition_Tables-20260817.md), the QCA preregistration (O2-P5_QCA_Preregistration-claude-opus-5-20260818.md), the truth-table report (O2-P5_QCA_TruthTable_Report-claude-opus-5-20260818.md), the existential judgment and its review (O2-P5_H0Existential_Judgment-gpt-5.6-sol-20260819.md and O2-P5_H0Existential_Judgment_Review-kimi-k3-20260819.md), the two motivating essays and the literature review, the permanent decision ledger (working-documents/Decision_Record.md, whose dispositions carry stable TTE-DEC identifiers), the shared handoff record (LLM_HANDOFF.md, which maintains the complete governed inventory), the structural verifier and its README, and the deterministic figure and build scripts. Revise this statement only after an authorized deposit is read-back verified.

SHA-256 and the structural verifier attest only pinned byte identity and structure, not truth, validation, or freedom from shared error. The visual script hash-gates its inputs and changes no plotted content.

**References**

Alexander, Jeffrey C., and Philip Smith, eds. 2005. *The Cambridge Companion to Durkheim*. Cambridge: Cambridge University Press.

Asad, Talal. 1993. *Genealogies of Religion: Discipline and Reasons of Power in Christianity and Islam*. Baltimore: Johns Hopkins University Press.

Atran, Scott. 2002. *In Gods We Trust: The Evolutionary Landscape of Religion*. Oxford: Oxford University Press.

Becker, Ernest. 1973. *The Denial of Death*. New York: Free Press.

Bellah, Robert N. 1967. "Civil Religion in America." *Daedalus* 96 (1): 1–21.

Berger, Peter L. 1967. *The Sacred Canopy: Elements of a Sociological Theory of Religion*. Garden City, NY: Doubleday.

Besant, Annie Wood, trans. 1922. *The Bhagavad-Gita, or The Lord's Song*. 4th ed. Madras: G. A. Natesan & Co.

Bostrom, Nick. 2003. "Are You Living in a Computer Simulation?" *The Philosophical Quarterly* 53 (211): 243–55.

Boyer, Pascal. 2001. *Religion Explained: The Evolutionary Origins of Religious Thought*. New York: Basic Books.

Camus, Albert. 1991a. *The Myth of Sisyphus and Other Essays*. Translated by Justin O'Brien. 1st Vintage International ed. New York: Vintage Books.

Camus, Albert. 1991b. *The Rebel: An Essay on Man in Revolt*. 1st Vintage International ed. New York: Vintage Books.

Chaves, Mark. 2010. "SSSR Presidential Address: Rain Dances in the Dry Season: Overcoming the Religious Congruence Fallacy." *Journal for the Scientific Study of Religion* 49 (1): 1–14.

Collins, Randall. 2004. *Interaction Ritual Chains*. Princeton, NJ: Princeton University Press.

Durkheim, Émile. (1893) 1984. *The Division of Labor in Society*. Translated by W. D. Halls. New York: Free Press.

Durkheim, Émile. (1897) 1951. *Suicide: A Study in Sociology*. Translated by John A. Spaulding and George Simpson. New York: Free Press.

Durkheim, Émile. (1912) 1995. *The Elementary Forms of Religious Life*. Translated by Karen E. Fields. New York: Free Press.

Foley, John. 2008. *Albert Camus: From the Absurd to Revolt*. Stocksfield: Acumen.

Fuller, Christopher J. 2004. *The Camphor Flame: Popular Hinduism and Society in India*. Rev. ed. Princeton, NJ: Princeton University Press.

Geertz, Clifford. 1973. *The Interpretation of Cultures*. New York: Basic Books.

Goff, Philip. 2019. *Galileo's Error: Foundations for a New Science of Consciousness*. New York: Pantheon Books.

Hackett, Conrad, Marcin Stonawski, Yunping Tong, Stephanie Kramer, Anne Shi, and Dalia Fahmy. 2025. *How the Global Religious Landscape Changed From 2010 to 2020*. Washington, DC: Pew Research Center. https://www.pewresearch.org/religion/2025/06/09/how-the-global-religious-landscape-changed-from-2010-to-2020/.

Lukes, Steven. 1985. *Emile Durkheim: His Life and Work: A Historical and Critical Study*. Stanford, CA: Stanford University Press.

Norenzayan, Ara. 2013. *Big Gods: How Religion Transformed Cooperation and Conflict*. Princeton, NJ: Princeton University Press.

Ragin, Charles C. 1987. *The Comparative Method: Moving Beyond Qualitative and Quantitative Strategies*. Berkeley: University of California Press.

Ragin, Charles C. 2008. *Redesigning Social Inquiry: Fuzzy Sets and Beyond*. Chicago: University of Chicago Press.

Rihoux, Benoît, and Charles C. Ragin, eds. 2009. *Configurational Comparative Methods: Qualitative Comparative Analysis (QCA) and Related Techniques*. Thousand Oaks, CA: SAGE.

Sagi, Avi. 2002. *Albert Camus and the Philosophy of the Absurd*. Translated by Batya Stein. Amsterdam: Rodopi.

Schneider, Carsten Q., and Claudius Wagemann. 2012. *Set-Theoretic Methods for the Social Sciences: A Guide to Qualitative Comparative Analysis*. Cambridge: Cambridge University Press.

Smith, Gregory A., Alan Cooperman, Becka A. Alper, Besheer Mohamed, Chip Rotolo, Patricia Tevington, Justin Nortey, Asta Kallo, Jeff Diamant, and Dalia Fahmy. 2025. *Decline of Christianity in the U.S. Has Slowed, May Have Leveled Off: Findings from the 2023–24 Religious Landscape Study*. Washington, DC: Pew Research Center. https://www.pewresearch.org/religion/2025/02/26/decline-of-christianity-in-the-us-has-slowed-may-have-leveled-off/.

Solomon, Sheldon, Jeff Greenberg, and Tom Pyszczynski. 2015. *The Worm at the Core: On the Role of Death in Life*. New York: Random House.

Sosis, Richard, and Candace Alcorta. 2003. “Signaling, Solidarity, and the Sacred: The Evolution of Religious Behavior.” *Evolutionary Anthropology* 12 (6): 264–74.

Stark, Rodney, and Roger Finke. 2000. *Acts of Faith: Explaining the Human Side of Religion*. Berkeley: University of California Press.

Strawson, Galen. 2006. “Realistic Monism: Why Physicalism Entails Panpsychism.” *Journal of Consciousness Studies* 13 (10–11): 3–31.

Sujato, Bhikkhu, trans. 2018a. *Linked Discourses: A Plain Translation of the Saṁyutta Nikāya*. SuttaCentral. https://suttacentral.net/sn22.59/en/sujato.

Sujato, Bhikkhu, trans. 2018b. *Long Discourses: A Faithful Translation of the Dīgha Nikāya*. SuttaCentral. https://suttacentral.net/dn2/en/sujato.

Sujato, Bhikkhu, trans. 2018c. *Middle Discourses: A Lucid Translation of the Majjhima Nikāya*. SuttaCentral. https://suttacentral.net/mn72/en/sujato.

Taylor, Charles. 2007. *A Secular Age*. Cambridge, MA: Belknap Press of Harvard University Press.

Ṭhānissaro Bhikkhu, trans. n.d. “To Vacchagotta on Fire (Aggi-vacchagotta Sutta, MN 72).” *dhammatalks.org*. Accessed August 18, 2026. https://www.dhammatalks.org/suttas/MN/MN72.html.

Whitehouse, Harvey. 2004. *Modes of Religiosity: A Cognitive Theory of Religious Transmission*. Walnut Creek, CA: AltaMira Press.

Scripture editions and other held sources not used directly in this synthesis remain documented in the coding record rather than being added as uncited works here; internal study records are inventoried under “Data, Materials, and Audit-Record Availability” rather than in this reference list.