

The Devil's in the Details—and So Is God: The Illogic of Dual-Agent Theism

Phil Stilwell

Independent Scholar

This paper analyzes the explanatory structure of theistic ideologies that posit both a benevolent agent (God) and malevolent counterparts (demons and morally depraved humans) as exhaustive causes of worldly events. By granting every occurrence to one of these categories, such frameworks achieve *explanatory closure* at the cost of *explanatory sterility*. No event can fall outside the ontology: prosperity is attributed to divine favor, misfortune to demonic interference, and moral evils to human depravity construed as Satan's instrumentality. While this produces the semblance of coherence, it collapses into tautology. Using both symbolic logic and Bayesian formalism, the paper demonstrates that dual-agent causality systems admit no falsifiers and generate Bayes factors approximating unity across outcomes, leaving them evidentially inert. The analysis situates this flaw within mainstream Christian categories, but the formal critique generalizes to any ideology that immunizes itself against disconfirmation by exhaustive agentic attribution. The conclusion is that such systems, far from supplying explanatory power, exemplify the illogic of ontological overreach.

Keywords: dual-agent theism, explanatory sterility, tautological coverage, epistemic elasticity, Bayesian analysis

Opening Narrative

Theistic ideologies frequently deploy a dualistic ontology in which the world is interpreted through the actions of opposed agencies: a benevolent God on one side and a malevolent counterforce on the other. In mainstream Christianity, this counterforce encompasses not only Satan and demons but also morally depraved humans acting under their influence (Moreland & Craig, 2003; Plantinga, 1974). The result is a framework in which all events—whether personal misfortunes, social catastrophes, or private moral failures—are already accounted for before they occur. An airplane crash may be ascribed to Satan’s destructive power; genocide may be explained as the consequence of human depravity in alignment with Satan; personal betrayal may be attributed to either direct demonic temptation or the sinful will of another. Whatever the event, it is guaranteed a place in the ontology.

This arrangement yields what can be termed *explanatory closure*. Nothing is left outside the system’s reach, and thus the ideology cannot encounter disconfirming evidence. Yet this closure simultaneously generates *explanatory sterility*. A hypothesis that predicts every possible outcome equally offers no discriminative power, and thus no genuine explanatory traction (Hume, 1748/2007; Sober, 2008). In what follows, the paper develops a conceptual lexicon, a formal logical and Bayesian framework, a vivid case study, and responses to objections, before situating the critique in broader epistemic contexts.

Lexicon of Key Terms

Ontic Dualism. The structuring of reality around two opposed supernatural agencies: a benevolent God and malevolent counterforces such as Satan, demons, or depraved humans. This framework ensures that every event is given an immediate agentic placement, thereby foreclosing neutral or non-agentic explanations (Moreland & Craig, 2003; Plantinga, 1974).

Dual-Agent Causality Trap. The logical snare produced when all events are captured by attribution either to God's providence or to malign agency. Because no possible outcome escapes these categories, the system is unfalsifiable. The "trap" lies in confusing coverage with explanatory power (Sober, 2008).

Explanatory Closure. The condition in which all conceivable outcomes are accounted for within the ontology of a worldview. While closure may appear to provide comprehensiveness, in practice it removes the possibility of evidential challenge: every datum is already anticipated, not by prediction, but by categorical absorption (Hume, 1748/2007).

Explanatory Sterility. The absence of evidential traction within a framework. A hypothesis is sterile when it predicts all possible outcomes equally well, yielding no differential confirmation or disconfirmation. In Bayesian terms, such a system produces Bayes factors near unity, no matter the observed evidence (Jaynes, 2003; Sober, 2008).

Tautological Coverage. Explaining events by reference to categories that merely restate the phenomenon: "If good, God; if evil, Satan or human sin." This is functionally tautological labeling rather than explanation (Hume, 1748/2007).

Epistemic Elasticity. The ability of an ideology to stretch and reclassify events so that no observation falls outside its scope, a hallmark both of conspiracy thinking and of dual-agent theism. Elasticity allows a system to appear robust under empirical challenge, but only by assimilating every counterexample into its structure (Sober, 2008).

Superficial Coherence. The appearance of internal consistency achieved by retrospective attributions that lack predictive bite (Sober, 2008).

Ad Hoc Subsumption. Reclassifying events after they occur to preserve consistency (e.g., “mystery” vs. “providence” vs. “temptation”), an ad hoc shield against falsification. Such auxiliaries preserve narrative harmony while eroding evidential content (Sober, 2008).

Possibility Laundering. Elevating speculative constructs (e.g., demonic causation) into assumed explanatory options without warrant, thereby laundering imagination into “live possibilities” (Jaynes, 2003).

Agentic Attribution. Assigning outcomes to invisible agents—divine, demonic, or depraved human—to supply intentional explanations, often at the expense of evidential discipline (Moreland & Craig, 2003; Plantinga, 2000).

Explanatory Vacuity. Explanations that cannot be wrong and thus cannot be informative. Vacuity maintains interpretive dominance while sacrificing evidential content (Hume, 1748/2007).

Bayesian Sterility. When $P(E \mid H) = 1$ across outcomes, the Bayes factor remains at unity; the hypothesis cannot update credence (Jaynes, 2003; Carrier, 2012).

Ideological Immunization. Neutralizing disconfirming evidence by reinterpretation (e.g., tests of faith, demonic interference, or inscrutable providence) so the theory is insulated from data (Plantinga, 2000; Swinburne, 2003).

Metaphysical Plugging. Filling explanatory gaps with unverifiable entities rather than reserving probability for the unknown, halting inquiry (Hume, 1748/2007).

Circular Ontology. Inferring the reality of God, Satan, and depravity from the very events they are invoked to explain; the ontology validates itself by its own categories (Carrier, 2014).

Explanatory Monopoly. Claiming exclusive interpretive authority over all outcomes. By monopolizing interpretation, the framework prevents genuine comparison and falsification (Sober, 2008).

Formalization: Dual-Agent Causality and Explanatory Sterility

Logical Skeleton

Define predicates: $O(e)$ = event e occurs; $G(e)$ = caused by God; $D_s(e)$ = caused by Satan/demons; $D_h(e)$ = caused by depraved humans.

$$\forall e [O(e) \rightarrow (G(e) \vee D_s(e) \vee D_h(e))].$$

Annotation. Every event is assigned to God, Satan, or depraved humans. No observation can yield $\neg(G(e) \vee D_s(e) \vee D_h(e))$. The framework is self-sealing (Hume, 1748/2007; Sober, 2008).

Bayesian Skeleton

Let $\mathcal{H} = \{H_G, H_{D_s}, H_{D_h}\}$ denote a partition: H_G = “God caused E ”; H_{D_s} = “Satan/demons caused E ”; H_{D_h} = “depraved humans caused E .” Standard Bayesian reasoning requires priors across all live hypotheses to sum to unity and to include reserves for unknown mechanisms:

$$\sum_i P(H_i) + P(H_{\text{unk}}) = 1.$$

Annotation. Dual-agent frameworks often set $P(H_{\text{unk}}) = 0$, enforcing closure at the cost of epistemic humility (Jaynes, 2003).

Mixture Likelihood. By the law of total probability:

$$P(E | I) = \pi_G P(E | H_G) + \pi_{D_s} P(E | H_{D_s}) + \pi_{D_h} P(E | H_{D_h}),$$

Annotation. With π . as priors. If each branch is flexibly tuned post hoc to “fit” E , the mixture becomes nearly flat across outcomes (Sober, 2008).

Specificity Penalty. Targeting a *specific* event pays a penalty:

$$P(H_{G,\text{specific}}) = P(H_{G,\text{generic}}) \times P(\text{this particular target}).$$

Annotation. Treating generic providence as if it implied targeted intervention inflates priors (Jaynes, 2003; Carrier, 2012).

Dependence Correction. Apparent multiple supports are often dependent. An effective-sample-size correction captures this:

$$n_{\text{eff}} \approx \frac{n}{1 + (n - 1)\rho},$$

Annotation. With average pairwise correlation ρ , testimony/interpretations are not independent; evidential force is overcounted (Jaynes, 2003; Carrier, 2012).

Likelihood Flatness & Bayes Neutrality. If every outcome can be narrated under any branch, then

$$P(E \mid H_G) \approx P(E \mid H_{D_s}) \approx P(E \mid H_{D_h}) \approx 1,$$

and the Bayes factor against a baseline H_0 is

$$BF_{I:H_0}(E) = \frac{P(E \mid I)}{P(E \mid H_0)} \approx 1.$$

Annotation. A theory that predicts everything equally predicts nothing; posterior odds barely shift (Jaynes, 2003; Sober, 2008).

Posterior Inertia. Posterior odds satisfy

$$\frac{P(I \mid E)}{P(H_0 \mid E)} = \frac{P(I)}{P(H_0)} \times BF_{I:H_0}(E).$$

Annotation. With $BF \approx 1$, posteriors mirror priors. Closure yields *Bayesian sterility* (Carrier, 2012; Sober, 2008).

Core Evidence and Analysis

The elasticity of dual-agent theism is evident in diverse domains. Natural disasters may be seen as divine judgment, demonic attack, or the byproduct of human sin; personal tragedies are assigned to providence, testing, or malevolent interference; atrocities such as genocide are explained as depravity while simultaneously framed as Satanic influence. Everyday ambiguities—illness, recovery, prosperity, misfortune—are all absorbed without risk. The result is not resilience but *vacuity*: nothing can count against the ontology (Hume, 1748/2007; Sober, 2008).

Case Study: Elastic Explanations in a Plane Crash

To illustrate elasticity, consider a commercial plane crash in which ninety percent of passengers perish. For an evidentially disciplined framework, such an event would test explanatory power: some outcomes should weigh against the hypothesis, others might support it. Within mainstream Christianity, however, because God, Satan, and depraved humans are treated as exhaustive causal agents, every outcome is assimilated into the ontology—a paradigm instance of *tautological coverage* (Carrier, 2012; Sober, 2008).

- **The spared traveler.** A passenger misses the flight due to a car accident. This is interpreted as providential sparing, an example of divine protection that “redirects” misfortune (Plantinga, 2000; Swinburne, 2003).
- **The last-minute transferee.** A traveler switches onto the doomed flight. If deceased, this is punishment, inscrutable timing, or a Satanic snare; if survived, it is testimony to providence.
- **The child saved by a dying parent.** The child’s survival is framed as direct providence; the parent’s death becomes a sanctified sacrifice permitted for a greater good.

- **The bereaved parent.** A surviving mother who loses her child is told this suffering is a divinely ordained test of faith, ultimately for sanctification.
- **The elderly couple.** Their joint death is read as merciful closure and the gift of finishing life “together in God’s timing.”
- **The newlyweds.** Their deaths are seen as bittersweet but providentially meaningful: a mysterious plan that dignifies their love through martyr-like witness.
- **The rival-religion adherent.** The death of a member of a rival faith is interpreted as just judgment or as divine permission following persistent rejection.
- **Post-crash conversion.** A survivor’s conversion is heralded as proof that God redeems tragedy, turning evil to good.
- **The scarred believer.** Severe disfigurement is cast as a refining trial that deepens faith and witness.
- **The unscathed believer.** Walking away unharmed is framed as a miraculous sign of protection and answered prayer.
- **The spared unbeliever.** Survival becomes “one more chance” to repent; continued unbelief is resistance to grace.
- **The deceased unbeliever.** Death is read as consequence of rejection or as an instance of God’s justice.
- **The heroic crew member.** Dying while saving others is orchestrated as a noble end that glorifies God through sacrificial love.
- **The pregnant victim.** The unborn child is said to be “called home early,” a mercy beyond human comprehension.

- **The atheist who prays.** Last-minute prayer is taken as implicit acknowledgment that God is real, validating theism.
- **The believer's peaceful death.** Dying in prayer is treated as confirmatory evidence of divine presence and "dying grace."
- **Split families.** Mixed outcomes within a family indicate distinct divine purposes tailored to each soul's sanctification and mission.
- **The late rescuer.** Arriving too late is read as humbling providence: a lesson in dependence on God's timing.
- **National response.** The country's reaction to the crash is interpreted as a call to collective repentance, a divinely permitted warning.
- **The executive cut short.** An executive's death is narrated as God's humbling of worldly ambition.
- **The pastor's death.** A pastor dying in the crash is a "martyr's witness," intended to strengthen others' faith.
- **The wealthy survivor.** A donor spared is preserved to fund kingdom work, a providential investment in ministry.
- **The childless couple.** Their deaths are mercifully framed as sparing future grief otherwise borne by children.
- **The survivor's apostasy.** If a survivor later leaves the faith, this is Satanic ensnarement in the aftermath.
- **The missionary survivor.** Turning to missions is credited as God converting tragedy into evangelistic opportunity.

Even the crash itself is folded into sovereignty (it serves hidden goods) or demonic attack (malevolent agency “permitted” within sovereignty). Either way, the event is *explained in advance*, with no risk (Swinburne, 2003; Plantinga, 2000). Formally,

$$\forall e [O(e) \rightarrow (G(e) \vee D_s(e) \vee D_h(e))]$$

ensures that the marginal likelihood already absorbs all outcomes; Bayes factors remain near unity. The result is not robustness but *explanatory sterility* (Jaynes, 2003; Carrier, 2012; Sober, 2008).

Counterarguments and Responses

Free Will Defense

Appeals to human freedom explain evil as the cost of agency (Plantinga, 1974). But within the triad (God, demons, depravity), free will becomes merely one branch of a partition that already covers all outcomes. The logical schema above shows why this is not a constraint but a re-description. In Bayesian terms, distributing probability mass among the three branches leaves $P(E | H) = 1$ across relevant outcomes; no discriminative power is introduced (Jaynes, 2003; Sober, 2008).

Appeal to Mystery

If apparently cruel or incoherent outcomes are answered by inscrutable purposes (Plantinga, 2000; Swinburne, 2003), then for all observed E , $P(E | H) = 1$. The Bayes factor is fixed near unity and the hypothesis ceases to interact with evidence (Carrier, 2012; Sober, 2008).

Comprehensiveness as Virtue

Claiming that “theism explains everything” confuses closure with explanation. Scientific models earn credit by risking refutation and by weighing outcomes differentially (Sober, 2008). If every event confirms H , then no event *counts* (Hume, 1748/2007).

Sophisticated Theological Systems

Plantinga's free-will defense, Molinist middle knowledge, and multiperspectival epistemologies refine mechanism, not evidential structure (Plantinga, 1974, 2000; Moreland & Craig, 2003). The triadic elasticity remains: all outcomes are pre-classified into God/Satan/depravity. Likelihoods remain effectively flat, preserving neutrality in the Bayes factor (Jaynes, 2003; Sober, 2008).

Discussion and Broader Implications

The same structure appears in conspiracy theories (absence of evidence becomes evidence of concealment) and in totalizing ideologies that attribute failures to sabotage but successes to inherent superiority (Sober, 2008). The broader lesson is methodological: *closure without falsifiability is sterility*. Progress requires risking disconfirmation and reserving probability for the unknown (Jaynes, 2003; Hume, 1748/2007).

The Missing Elements of Rational Scrutiny

If dual-agent theism is explanatorily sterile, why does it retain plausibility? Several *absent safeguards* and *systematic errors* help explain its persistence.

Absence of Probability Reserves

Sound inference preserves *probability mass for the unknown*. Forcing immediate assignment to God/Satan/depravity eliminates uncertainty and prevents learning. Formally, refusing a live H_{unk} collapses posterior dynamics: with no capacity to absorb surprises, posteriors cannot move (Jaynes, 2003).

Neglect of Base Rates

When background frequencies are ignored, rare agentic explanations are overweighted. Proper conditioning on mundane base rates (mechanical failure, human error, stochastic risk) would typically dominate; elasticity masks this by selective narration (Sober, 2008).

Suppression of Falsifiability

A theory gains content by ruling things out (Hume, 1748/2007). Where no outcome disconfirms theism because it is reallocated to a branch *ex post*, content drops to near zero. This is the essence of *tautological coverage* (Sober, 2008).

Ad Hoc Elasticity and Complexity Penalties

Each unexpected outcome prompts a fresh auxiliary (“test,” “discipline,” “mystery”). These ad hoc add-ons increase description length without increasing predictive success. In Bayesian model comparison, marginal likelihood penalizes such complexity, lowering posterior support (Jaynes, 2003).

Confirmation by Anecdote and Selective Likelihoods

Anecdotes are promoted to decisive evidence while disconfirming cases are reinterpreted. Selective conditioning (“count the hits, explain the misses”) destroys calibration and simulates evidential accumulation where none exists (Carrier, 2012, 2014).

Ambiguity Laundering

Ambiguous data are retrofitted into preferred categories after the fact, inflating perceived fit and collapsing predictive testing. Post hoc classification eliminates the very asymmetry evidence must have to inform belief (Hume, 1748/2007).

Prior Dogmatism and Posterior Inertia

If priors on theism/demonic agency are set dogmatically high, posterior inertia follows regardless of data. Rational priors must be elastic enough to move under differential likelihoods; otherwise $P(H | E) \approx P(H)$ by design (Jaynes, 2003).

Asymmetric Explanatory Standards

Natural explanations are rejected for being incomplete, while theistic explanations are accepted despite being unfalsifiable. This reverses methodological burden: the more

testable the claim, the harder it is treated; the less testable, the more easily it is accepted (Sober, 2008).

Teleological Overreach and Global Purpose Attribution

Global-purpose narratives supply meaning but erase evidential constraints. Assigning teleology to every event ensures that no observation is “out of model,” securing comfort at the cost of information (Plantinga, 2000; Swinburne, 2003).

Failure of Comparative Evaluation

Rival hypotheses (mechanical failure, human error, stochastic risk) are not compared on common likelihoods. Without a comparative metric, monopolistic interpretation is guaranteed; the dual-agent model “wins” by excluding rivals from the scoreboard (Sober, 2008).

Dependence and Echo Chambers

Communal interpretations interdepend, yielding correlated testimonies that masquerade as multiple independent confirmations. Effective-sample-size corrections reveal far less evidential force than narrative suggests (Jaynes, 2003; Carrier, 2012).

Conclusion

By positing God, Satan, and depraved humans as exhaustive causal agents, dual-agent theism achieves closure but collapses into sterility. The lexicon and formalization showed that no observation can falsify it, while the Bayesian analysis showed that it produces no evidential updates. What appears as comprehensiveness is revealed as *tautological coverage*: predicting every outcome equally and therefore explaining none (Hume, 1748/2007; Jaynes, 2003; Sober, 2008; Carrier, 2012).

The plane-crash case study illustrated how contradictory fates are seamlessly assimilated, ensuring that no result counts against the system. Appeals to free will, mystery, or theological sophistication refine narrative mechanics but do not restore

evidential constraint (Plantinga, 1974, 2000; Swinburne, 2003; Moreland & Craig, 2003).

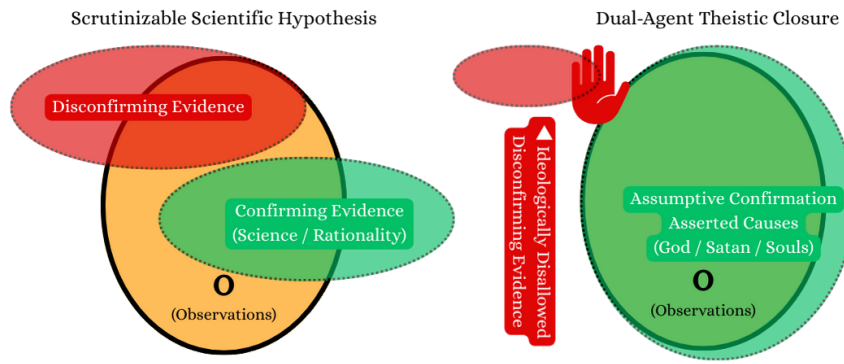
The broader moral is epistemic: systems that treat every outcome as confirmatory confuse closure with explanation. The same sterility appears in conspiracy thinking and totalizing ideologies. Inquiry requires probability reserves, testable commitments, and the humility to let evidence bite (Hume, 1748/2007; Jaynes, 2003; Sober, 2008).

References

- Carrier, R. (2012). *Proving history: Bayes's theorem and the quest for the historical Jesus*. Prometheus Books.
- Carrier, R. (2014). *On the historicity of Jesus: Why we might have reason for doubt*. Sheffield Phoenix Press.
- Hume, D. (2007). *An enquiry concerning human understanding* (P. Millican, Ed.). Oxford University Press. (Original work published 1748)
- Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge University Press.
- Moreland, J. P., & Craig, W. L. (2003). *Philosophical foundations for a Christian worldview*. IVP Academic.
- Plantinga, A. (1974). *God, freedom, and evil*. Eerdmans.
- Plantinga, A. (2000). *Warranted Christian belief*. Oxford University Press.
- Sober, E. (2008). *Evidence and evolution: The logic behind the science*. Cambridge University Press.
- Swinburne, R. (2003). *The resurrection of God incarnate*. Oxford University Press.
- Stilwell, P. (2024). *Bayesian pitfalls in resurrection apologetics*. Unpublished manuscript.

Appendix

Visual Contrast of Explanatory Frameworks

**Figure A1**

Contrast between a scrutinizable scientific model (left) and a dual-agent theistic closure (right).

Explanatory Note. The left side depicts a scientific or rational framework in which observations (O) may either confirm or disconfirm a hypothesis. Some data fall into $O \cap P(H)$ and thus support the hypothesis, while other data fall into $F(H) = O \cap \neg P(H)$ and function as falsifiers. This structure embodies the principle of falsifiability and preserves explanatory discrimination.

The right side depicts a dual-agent theistic framework in which all observations are absorbed into an overarching ontology of God, Satan, or human depravity. Because disconfirming evidence is excluded *a priori*, every observation is treated as confirmatory. This yields *explanatory sterility*: the framework appears comprehensive but is immunized against falsification, collapsing into tautology. The figure visually underscores the epistemic gap between testable scientific models and the self-sealing nature of dual-agent theism.