

Explanatory Vacuity

Dual-Agent Theism, Tautological Coverage, and the Failure of Risky Explanation

Phil Stilwell

Independent Scholar

ORCID: 0009-0009-3725-8682

Revised August 2026

Central claim

A dual-agent scheme does not gain public confirmation merely because it can assign favorable outcomes to a deity and unfavorable outcomes to an adversary, sin, discipline, testing, or mystery. Unless those categories yield independently specified likelihood differences, exclude outcomes, track complete denominators, and permit belief revision, they provide interpretive coverage rather than risky explanation. This conclusion does not disprove spiritual agents or deny the pastoral value of religious meaning.

Preview: Coverage is not discrimination

Starting claim	Required risk	Actual pressure	Rational conclusion
Every outcome receives a coherent spiritual interpretation.	The scheme must make some outcomes more expected than relevant alternatives before the result is known.	Opposite results can be reassigned to a different agent, purpose, or mystery without lowering confidence.	Successful assignment shows coverage; confirmation requires advance discrimination, denominators, and revision.

The target is the use of unrestricted spiritual assignment as public evidence. A constrained religious explanation may remain evidentially serious, and a pastoral interpretation may remain meaningful, without claiming that the event confirms a worldview.

Visual preview:

Explaining Everything, Explaining Nothing: The Risk of Explanatory Vacuity

THE PROBLEM: THE “EVERYTHING” TRAP

Tautological Coverage

When a system's categories are so broad that no possible observation falls outside them.

Providence

The Plane-Crash Test

Mystery

The framework simply relabels reality

The Failure of Riskless Explanation

An explanation that cannot be disappointed by any outcome provides no real information.

THE SOLUTION: RISK-BEARING EXPLANATION

Constrain Expectations



Genuine explanations must specify what would count against the model before evidence arrives.

Honor Ordinary Causes



Prioritize mechanical, biological, and historical mechanisms over spiritual relabeling.



Emotional Comfort
(Meaning)



Objective Proof
(Evidence)

Distinguish Meaning from Evidence

Separate emotional comfort and pastoral meaning from claims of objective proof.

Read the image from left to right. The left panel depicts the self-sealing rule: a favorable result is assigned to divine action, while an unfavorable result can be reassigned to an adversary, human fault, discipline, testing, or mystery. The plane-crash examples show why opposite outcomes and uncounted cases must remain in the same evidential denominator. The right panel gives the remedy: state expectations before the outcome, preserve ordinary causal investigation, count successes and failures together, and reduce confidence when specified expectations fail. The image diagnoses an evidential method; it does not assert that spiritual agents are unreal or that religious meaning is illegitimate.

Abstract

Some religious interpretations assign favorable events to a deity and adverse events to an adversary, human sin, discipline, testing, judgment, or mysterious providence. Because every outcome receives a spiritual label, the scheme can appear unusually comprehensive. This paper argues that comprehensive assignment is not yet public explanation. The relevant question is whether the scheme, before the outcome is known, makes the observed result more expected than a relevant ordinary alternative and accepts evidential loss when contrary results occur.

The analysis separates four functions often compressed into the word *explain*: assigning meaning, locating a cause, guiding intervention, and confirming a worldview. It then develops a likelihood-based symmetry test and applies it to a balanced plane-crash ledger. Survival and death, conversion and unbelief, a missed flight and a last-minute boarding can each be assimilated by changing the spiritual label. That flexibility does not prove that the labels are false; it shows that the successful assignment alone supplies no independently specified likelihood advantage. Selective testimony, hidden denominators, and unrestricted auxiliary explanations further prevent learning.

The conclusion is bounded. The argument neither disproves a deity, an adversarial spirit, or providence nor denies that religious narratives can console, orient action, or express communal commitments. It shows that an unrestricted dual-agent scheme forfeits public confirmation from the very outcomes it automatically absorbs. A risk-bearing religious explanation must specify comparative expectations in advance, retain complete outcome records, cooperate with ordinary causal inquiry, and permit confidence to decrease. Without those disciplines, universal coverage is interpretation rather than evidential discrimination.

Keywords: explanatory vacuity; dual-agent theism; providence; spiritual warfare; likelihood; confirmation; falsifiability; base rates; public evidence

1 The Attraction and the Mistake

A framework that never runs out of answers can feel explanatorily powerful. A recovery is a blessing. A failure is discipline. A temptation is attributed to an adversarial spirit. An unanswered prayer expresses wiser providence. A survivor's conversion reveals purpose; a survivor's unbelief reveals resistance. The framework appears to cover both ordinary life and exceptional tragedy without remainder.

The mistake is to move directly from *can be interpreted within the framework* to *confirms the framework*. Compatibility is cheap. Many rival accounts can accommodate an event after it occurs. Public confirmation requires discrimination: the observation must be more expected under the proposed account than under a relevant alternative, given independently supported assumptions. Bayesian accounts make this comparative structure explicit, while inference-to-the-best-explanation approaches likewise assess explanatory success against rivals rather than in isolation (Howson and Urbach 2006; Lipton 2004).

The target of this paper is therefore narrow. A *dual-agent scheme* here is an interpretive system that assigns favorable outcomes primarily to a benevolent deity and adverse outcomes to one or more flexible auxiliaries: an adversarial spirit, human sin, discipline, testing, judgment, hidden goods, or mystery. The target is not every belief in providence or spiritual conflict. It is the public evidential use of a scheme whose assignments are chosen after the event and whose confidence is insulated from opposite outcomes.

The relevant alternative is also local rather than metaphysically exhaustive. Call it the *ordinary causal account*: the event is explained through known or investigable mechanisms, base rates, human choices, institutions, chance processes, and acknowledged uncertainty. This account need not deny all spiritual reality. It simply asks what the public record supports before a spiritual interpretation is added.

The governing distinction

If opposite outcomes can be absorbed with equal ease by changing the spiritual label after the fact, then the scheme's successful assignment of the actual outcome shows interpretive coverage. It does not, without further likelihood constraints, show that the outcome publicly confirms the scheme over ordinary causes.

This distinction does not accuse believers of insincerity. Grief and gratitude naturally invite narrative meaning. The problem arises only when a narrative that helps a community inhabit an event is treated as new evidence for the worldview that supplied the narrative.

2 Four Jobs Hidden by “Explain”

The word *explanation* can name several different achievements. Keeping them separate prevents a useful religious practice from being credited with a different evidential achievement it has not earned.

2.1 Meaning assignment

Meaning assignment places an event within a story about identity, value, or purpose. A community may understand survival as an obligation to serve, suffering as an occasion for solidarity, or grief as something carried in hope. Such interpretations can alter conduct and sustain people without identifying a hidden efficient cause. Their standard is practical or existential adequacy: does the story orient life truthfully and humanely?

2.2 Causal location

A causal explanation identifies factors that produced the event and shows how changing them would alter the result. Contemporary interventionist approaches emphasize this difference-making role (Woodward 2003). In an aviation disaster, investigators examine weather, maintenance, design, training, communication, and human decisions. The explanation earns force by connecting the record to mechanisms and counterfactual dependence, not by attaching a morally resonant label.

2.3 Expectation and prediction

An explanatory model can also organize expectation. Hempel's influential account tied explanation to law-governed expectability, even though later work has rejected the idea that every explanation has the same deductive form (Hempel 1965). The durable lesson is modest: a successful empirical explanation constrains what should be expected. It does not merely announce, after the outcome, that the result falls somewhere inside an unlimited interpretive vocabulary.

2.4 Worldview confirmation

Finally, an observation may be used to raise confidence in a broader worldview. This is the strongest job and requires a comparator. Swinburne's theistic cumulative case, for example, explicitly treats evidence as confirmatory when it is expected on the proposed hypothesis and not otherwise expected, alongside simplicity and fit with background knowledge (Swinburne 2004). Dawes, while much more skeptical of theistic explanation, similarly grants that agent-based religious explanation is possible in principle and asks whether it meets ordinary explanatory standards (Dawes 2009). The dispute is therefore not between religion and a rule invented to exclude it. The shared question is whether the proposed explanation earns comparative advantage.

One event can succeed at one job and fail at another. A funeral narrative may assign humane meaning while remaining silent about mechanical cause. A causal report may identify a failed component

while saying nothing about ultimate purpose. A devotional interpretation may express an existing commitment without becoming independent evidence for that commitment. Clarity about the job prevents category inflation.

3 Comparative Likelihood and Risky Expectation

Suppose an observation is denoted by O , a dual-agent hypothesis by H , and an ordinary alternative by N . The evidential question is not whether H can accommodate O . It is whether O was more expected under H than under N . A likelihood ratio expresses that comparison:

$$\Lambda(O; H, N) = \frac{P(O | H)}{P(O | N)}. \quad (1)$$

In plain language, the numerator asks how expected the observation is if the dual-agent hypothesis is true; the denominator asks how expected it is under the ordinary comparison. A ratio above one favors H on this item, a ratio of one is neutral, and a ratio below one favors N . This statement concerns the direction of evidence, not the final probability of either hypothesis, which also depends on prior information and other evidence (Howson and Urbach 2006).

The formula does not manufacture numbers. Both likelihoods need reasons. A defender must specify what features of the agents, intentions, practices, or circumstances make one outcome more expected. Those assumptions must be available independently of simply redescribing the outcome after it occurs. If no likelihoods can be responsibly estimated even comparatively, humility is required: the observation's evidential direction remains unestablished.

Agency does not remove this burden. Human agents are not mechanical laws, yet intentions, character, capacities, opportunities, and costs support expectations about their actions. A hypothesis about an unconstrained agent whose reasons can fit any behavior will accommodate more data, but that very flexibility makes a particular behavior less discriminating. A personal explanation succeeds when evidence about the agent and setting narrows what should be expected.

Risk is the practical face of discrimination. A model takes evidential risk when it identifies in advance observations that would be less expected under it and would lower rational confidence. Popper made exposure to possible refutation central to empirical testing, although falsifiability is not a complete theory of meaning or confirmation (Popper 1959). The present argument needs only the weaker point: surviving a test supports a claim only if the claim could have fared worse.

This yields three different states that should not be conflated:

- **Specified inequality:** independent reasons support $P(O | H) > P(O | N)$, so O favors H on that comparison.
- **Specified equality:** reasons support equal likelihoods, so O is neutral between them.
- **Unspecified likelihoods:** the scheme can label O but supplies no disciplined comparison. No evidential advantage has been established.

Explanatory vacuity normally occupies the third state. The criticism is not that a flexible scheme must assign mathematically equal probabilities to every outcome. It is that post-hoc coverage, by itself, does not supply the likelihood asymmetry required for confirmation.

4 How Unrestricted Auxiliaries Erase Learning

Serious theories rarely confront evidence alone. They work with auxiliary assumptions about instruments, initial conditions, background mechanisms, or human choices. Contrary evidence can rationally lead investigators to revise an auxiliary rather than abandon the central claim. Lakatos therefore assessed

research programmes over time: protective adjustments may be progressive when they generate novel understanding, or degenerating when they only shelter the core from recalcitrant facts (Lakatos 1970).

The problem is not the existence of auxiliaries but their discipline. In an unrestricted dual-agent scheme, the auxiliary menu can assign any result to a different source:

- favorable result: blessing, protection, answer, favor, or providence;
- adverse result: adversarial attack, human sin, discipline, judgment, testing, hidden good, or mystery;
- no observable difference: hidden action, respect for ordinary processes, or reasons beyond human access;
- apparent counterexample: misinterpretation, insufficient faith, spiritual deception, or a deeper plan.

Each item may be logically possible. Logical possibility is not the same as evidential warrant. When the choice among auxiliaries is made only after the outcome, and no assignment carries a predictable cost for the core framework, the scheme records events without learning from them.

Boudry and Braeckman describe *immunizing strategies*: moves that prevent evidence from threatening a favored claim (Boudry and Braeckman 2011). Immunization can sometimes be warranted. An anomalous laboratory result may reflect a contaminated sample; an apparently broken promise may have an explanation. The warning sign is systematic asymmetry. Evidence favorable to the core is accepted at face value, while unfavorable evidence activates a costless auxiliary. Confidence ratchets upward but is not permitted to move downward.

The one-way update test

Ask whether the same evidential standard can both raise and lower confidence. If a success is counted as confirmation but a failure is always reassigned to an adversary, hidden purpose, or mystery, the framework has a one-way update rule. It is protected from loss, so accumulated “wins” cannot be read as ordinary confirmation.

The ratchet can be psychologically invisible because the auxiliary labels have moral and scriptural resonance. Yet resonance does not determine evidential direction. A disciplined use would specify when each auxiliary applies, what independent evidence supports it, and what observations would make that assignment less credible.

5 The Balanced Plane-Crash Ledger

A disaster brings the symmetry into view because many opposed outcomes occur inside one event. Consider a flight that crashes after a chain of ordinary failures. One traveler misses it because of a broken-down car; another boards at the last minute. Some passengers survive; others die. One survivor converts; another loses faith. Families report warnings, prayers, coincidences, guilt, gratitude, and anger.

The following ledger is intentionally balanced. It does not mock the interpretations or claim access to hidden purposes. It asks what happens evidentially when both an outcome and its opposite have readily available spiritual assignments.

Paired result	Available label for one result	Available label for the opposite	Evidential consequence
Missed flight / late boarding	Providence prevented the trip	A hidden purpose permitted or directed the trip	Travel timing alone does not select between the scheme and ordinary contingency
Survival / death	Protection, answered prayer, or unfinished purpose	Higher timing, judgment, testing, sacrificial witness, or mystery	Counting only survivors hides the full denominator; counting both preserves no obvious likelihood advantage
Recovery / lasting injury	Healing or special mercy	Discipline, testimony through suffering, adversarial damage, or mystery	Opposed medical outcomes remain assimilable unless differential rates and mechanisms were specified
Conversion / unbelief	Tragedy was used to awaken faith	A heart resisted, an adversary deceived, or grief became a test	Either psychological trajectory can be narrated as spiritually significant after it occurs
Prior warning / no warning	A dream, intuition, or delay was guidance	Silence itself required trust or protected freedom	Salient warnings need a base rate that includes forgotten impressions and uneventful trips
Community renewal / fracture	Suffering produced redemptive unity	Conflict exposed sin or intensified spiritual attack	Opposite social effects can each be assigned a role without risking the core claim

Three lessons follow. First, assignment is retrospective. The same label is not applied by a stable rule across cases; the vocabulary is selected to fit the result. Second, the actual outcome is compared with no counterfactual expectation. We are rarely told, before the flight, which distribution of survival, prayer reports, or conversions the hypothesis makes more likely. Third, attention is selective. A dramatic near miss becomes testimony, while thousands of similar delays on safe flights and unremembered intuitions disappear from view.

The test can be run prospectively. Before learning the record, ask a defender to state which possible distribution would favor the dual-agent scheme over ordinary causes, which would be neutral, and which would count against it. Require the same standards after the record is revealed. If no conceivable distribution would lower confidence, then the event cannot provide the upward evidential movement later claimed for it.

This conclusion is compatible with compassion. A survivor may experience gratitude and adopt obligations because of the event. A mourner may find a ritual vocabulary that makes continued life possible. Those are genuine human consequences. The evidential question is narrower: does the distribution publicly discriminate a spiritual scheme from the ordinary causal account? The ledger shows why meaning alone does not answer it.

6 Base Rates, Mechanisms, and Denominators

Extraordinary stories become persuasive partly by isolating a numerator. We hear of the person whose delay prevented boarding, not the many travelers whose delays were merely frustrating. We hear of a prayed-for recovery, not the full set of comparable patients, treatments, prayers, non-prayers, improvements, deteriorations, and reporting choices. Without the denominator, an event's rarity and comparative fit cannot be assessed.

A *base rate* is the frequency of an outcome in the relevant comparison class before the special interpretation is added. It matters because ordinary mechanisms already generate coincidences, recoveries, missed connections, vivid dreams, changes of belief, and survivor narratives. The more opportunities there are for a match, the less surprising one selected match becomes. The correct denominator must also include failed, forgotten, ambiguously described, and unpublished cases.

Mechanisms matter for a different reason. They connect explanation to control. Aviation investigation can identify a component, decision, or regulation whose alteration changes risk. Clinical inquiry can compare treatments and natural histories. Institutional inquiry can expose incentives and failures. These accounts may remain incomplete, but their errors are corrigible: additional evidence can distinguish rival mechanisms and improve intervention.

A spiritual layer need not compete with that work. A believer may hold that ordinary causes operate within a larger providential order. The evidential problem begins when the larger layer is treated as confirmed by whichever ordinary result occurs. Layering causes does not create discrimination. If the spiritual claim leaves all mechanical distributions unchanged and supplies no further risky expectation, the mechanical evidence cannot select it over a worldview that accepts the same mechanisms without that layer.

This point has ethical weight. Premature spiritual labeling can redirect attention away from preventable causes. An abusive institution may describe exposure as spiritual attack. A disease may be framed as judgment when public-health action is needed. Depression may be attributed solely to an adversarial spirit when clinical care is available. The claim is not that religious practices are always harmful or that causal inquiry is morally complete. It is that a label should not block investigation, accountability, or intervention.

7 The Conspiracy Comparison, Carefully Bounded

Conspiracy theories provide a useful structural comparison because hidden agency can explain both evidence and its apparent absence. A supporting document confirms the plot; a contrary document is planted. Expert agreement shows coordination; disagreement shows controlled opposition. Keeley analyzes how some conspiracy accounts expand skepticism toward the institutions that would ordinarily test them (Keeley 1999).

The comparison is diagnostic, not an equation between religious believers and conspiracists. Many conspiracies are real, many religious claims are argued with evidential discipline, and people in both domains vary widely in intellectual care. The shared danger is a self-sealing inference pattern:

1. posit a hidden agent capable of producing the observation;
2. treat favorable appearances as direct evidence;
3. reinterpret unfavorable appearances as evidence of deeper hidden activity;
4. preserve the core claim regardless of which appearance occurs.

Hidden agency is not inherently defective. Courts, historians, and investigators routinely infer concealed intentions. They succeed by connecting claims to restricted capacities, documented opportunities, stable motives, independent traces, and risky predictions. A proposed cover-up can be weakened when expected records, communications, or behavioral patterns fail to appear. Likewise, a religious agent explanation can in principle be constrained by a developed account of intentions and observable consequences (Dawes 2009).

The defect is unrestricted elasticity. If every missing trace is itself proof of concealment, investigation cannot discriminate. If every adverse religious outcome is proof of adversarial intensity or hidden benevolence, experience cannot correct the assignment. In both cases, the remedy is not a ban on hidden agents. It is a demand for independently supported constraints and symmetric updating.

8 Strongest Replies

8.1 Personal agency is not statistical regularity

A defender may object that persons respond differently in different circumstances. A deity is not a vending machine, and prayer is not a laboratory lever. This is correct. The likelihood approach does not require deterministic repetition or simplistic frequency claims. It requires some reason why the agent hypothesis makes the observed pattern more expected than its alternatives.

We routinely reason about agents through purposes, character, knowledge, opportunity, and cost. If every action and omission can equally express perfect wisdom, however, the character attribution stops narrowing expectation. Complexity can justify uncertainty; it cannot turn lack of discrimination into confirmation. The appropriate conclusion may be that the event is evidentially opaque.

8.2 Scripture independently supplies the categories

Another reply says that blessing, adversarial action, discipline, testing, and mystery are not improvised labels; they are supplied by an authoritative tradition. This may establish that the categories are internally available. It does not show which category applies to a particular event, how the assignments compare with ordinary alternatives, or why the event independently confirms the authority of the tradition.

If scriptural authority is assumed, the event can be interpreted within that framework. If the event is then offered as public evidence for the authority assumed in interpreting it, the reasoning becomes circular. A noncircular case needs independent criteria linking particular observations to particular categories and allowing failed expectations to count.

8.3 Spiritual and ordinary causes operate at different levels

A defender may accept every mechanical explanation while adding that a deity works through those mechanisms. The levels need not compete. This layered view avoids the practical danger of rejecting ordinary causes and may be metaphysically coherent.

Yet compatibility at another level is not comparative evidence. If the added spiritual claim predicts exactly the same publicly observable record as the ordinary account, the record is neutral between them. To gain evidential force, the higher-level claim must contribute some independently motivated pattern, constraint, or explanatory virtue beyond its ability to redescribe the completed causal chain.

8.4 The interpretation is pastoral, not evidential

This reply concedes the central distinction and may be entirely reasonable. A community can use inherited language to mourn, express hope, or identify obligations without claiming to have discovered a hidden cause. The paper does not require pastoral speech to function like a controlled study.

The boundary must then be kept. The same testimony cannot be exempted from evidential standards as pastoral expression and later counted as public confirmation. Speakers should say whether they are offering meaning, a causal claim, or evidence for a worldview. Different purposes deserve different assessments.

8.5 Conversions, healings, and fulfilled impressions show fruit

Apparent fruit can matter, but it requires a complete record and a comparator. Conversions after tragedy must be counted alongside unbelief, deconversion, and change toward other traditions. Healing testimony must be assessed alongside prognosis, treatment, spontaneous recovery, non-recovery, and reporting selection. Fulfilled impressions must be compared with vague, forgotten, revised, and unfulfilled impressions.

The point resembles Draper's use of comparative likelihood in the evidential problem of suffering: a distribution matters through how well it fits competing hypotheses, not merely because either can

accommodate it (Draper 1989). This paper does not adopt Draper's larger conclusion. It adopts the methodological demand that the complete pattern, rather than a selected success, carry the inference.

9 A Protocol for Risk-Bearing Religious Explanation

A religious explanation need not become a simple prediction machine to avoid vacuity. It does need practices that let observations constrain belief. The following protocol turns that requirement into an audit.

1. **Name the precise hypothesis.** State which agent, intention, capacity, and circumstance are being proposed. "Providence" by itself is too broad to generate a comparison.
2. **Separate the explanatory job.** Identify whether the claim assigns meaning, locates a cause, guides action, or offers public confirmation. Do not transfer success from one job to another without argument.
3. **Fix expectations before inspection.** State which observations are more expected, less expected, or neutral under the hypothesis compared with a named alternative. Avoid retrospective sharpening.
4. **Record the whole denominator.** Preserve successes, failures, ambiguous cases, non-events, opposite outcomes, and the rules used to classify them. Do not count only memorable testimony.
5. **Specify auxiliary rules.** Explain when adversarial action, discipline, testing, human fault, or mystery is licensed and what independent evidence distinguishes those assignments.
6. **Retain ordinary mechanisms.** Investigate base rates, causal pathways, treatment effects, institutional responsibility, and preventable factors before adding a spiritual layer.
7. **Permit downward revision.** State what result would reduce confidence and by how much in qualitative terms if numbers are unavailable. Apply the same evidential standard to success and failure.
8. **Preserve the unknown.** Where likelihoods, causes, or purposes cannot be responsibly compared, say so. Uncertainty is more informative than a costless label.

The protocol does not guarantee agreement. Likelihood judgments can remain contested, background assumptions can differ, and personal evidence may rationally affect participants differently from public observers. Its value is procedural: it reveals where the evidential work occurs and prevents universal assignability from masquerading as support.

The protocol also permits strong cases in principle. Suppose a community states a specific expectation in advance, defines a comparison class, records every relevant case, identifies a pattern that departs from ordinary causes, and revises confidence when replications fail. That practice would bear evidential risk. Whether the result ultimately supports a religious explanation would depend on the actual record, but the method would no longer be vacuous by construction.

10 Meaning Without Evidential Overreach

Human beings do not encounter tragedy only as statisticians. They grieve, remember, promise, worship, protest, and rebuild. A framework may help a community decide how to live after an event even when it does not identify why the event occurred. This practical orientation can be evaluated for honesty, compassion, moral consequence, and fit with the community's commitments.

Restraint protects that meaning. When hidden purpose is announced too quickly, sufferers can become props in someone else's theological narrative. A death becomes a lesson, an injury a test, and

survival a badge of favor before those most affected have spoken. Admitting ignorance can express moral seriousness rather than spiritual indifference.

The distinction also protects inquiry. A believer can say, “We will respond with care and hope,” while investigators say, “The crash followed from these failures.” Neither statement needs to displace the other. The first is a commitment about response; the second is a corrigible causal account. A third claim—“the crash publicly confirms our account of hidden agency”—carries a different burden and must face comparative likelihoods, denominators, and possible defeat.

Meaning therefore need not disappear when evidential overreach is removed. It becomes more accurately placed. Religious language can console, bind communities, and express values without converting every outcome into proof. Public explanation can remain answerable to mechanisms and evidence without pretending to settle ultimate purpose.

11 Conclusion

A dual-agent scheme can appear comprehensive because no event falls outside its vocabulary. Favorable outcomes can be assigned to a benevolent deity; unfavorable outcomes can be reassigned to an adversary, human sin, discipline, testing, judgment, hidden goods, or mystery. But complete assignability is coverage, not yet confirmation.

Public confirmation requires a disciplined comparison. The observed result must be more expected under the proposed hypothesis than under a relevant alternative for reasons specified independently of the result. Complete denominators must include opposed and failed cases. Ordinary mechanisms must remain visible. Most importantly, the framework must allow evidence to lower as well as raise confidence.

The plane-crash ledger exposes the problem without claiming knowledge of hidden reality. Survival and death, missed travel and last-minute boarding, conversion and unbelief can all receive coherent meanings. The successful existence of a meaning does not tell us which worldview the distribution favors. Where no likelihood asymmetry is supplied, the evidential direction remains unestablished.

The conclusion is therefore direct and bounded. Explanatory vacuity does not disprove a deity, an adversarial spirit, or providence. It does not deny the value of pastoral narratives. It shows that an unrestricted dual-agent scheme cannot claim public confirmation from outcomes it automatically absorbs. Risk-bearing explanation constrains expectation, honors ordinary causes, records the denominator, and permits revision. Without those disciplines, explaining everything becomes another way of learning nothing from what occurs.

A Formal Reconstruction of the Central Thesis

The reconstruction states the paper’s limited conclusion: successful post-hoc assignment does not by itself establish public confirmation. It does not conclude that the dual-agent hypothesis is false.

Syllogistic form

1. **Premise 1.** If the complete record publicly confirms the dual-agent hypothesis over the ordinary alternative, it supplies independently specified likelihood discrimination.
2. **Premise 2.** Unrestricted post-hoc absorption prevents the record’s successful assignment from supplying that discrimination.
3. **Premise 3.** The targeted scheme has unrestricted post-hoc absorption.
4. **Conclusion.** Therefore its successful assignment of the record does not establish public confirmation.

Symbols

H : the dual-agent hypothesis and its spiritual assignment rules.

N : the ordinary causal alternative using mechanisms, base rates, human choices, and uncertainty.

O : the complete record, including favorable, unfavorable, ambiguous, and unreported-opportunity classes.

$P(O | H)$: the likelihood of the complete record under H .

$P(O | N)$: the likelihood of the same complete record under N .

$\Lambda(O; H, N)$: the likelihood ratio comparing those expectations.

$D(H, N, O)$: independent reasons support a likelihood advantage for H over N without merely redescribing O .

$U(H, O)$: H absorbs O and relevant contrary outcomes post hoc by changing auxiliaries without evidential cost.

$C(H, N, O)$: the claim that O has been shown to publicly confirm H over N .

Every symbol defined above is used below. “Independently specified” does not mean theory-free; it means that the likelihood reason is not created solely by selecting a fitting label after observing the result.

Symbolic form

Likelihood diagnostic.

$$\Lambda(O; H, N) = \frac{P(O | H)}{P(O | N)}. \quad (\text{A1})$$

Annotation. Equation (A1) compares the same complete record under the two accounts. If justified likelihoods yield $\Lambda(O; H, N) > 1$, the record favors H on this comparison; if the ratio is one, it is neutral; if below one, it favors N . When the likelihoods are unsupported, the equation licenses no verdict.

$$D(H, N, O) \Rightarrow \Lambda(O; H, N) > 1. \quad (\text{A2})$$

Annotation. Equation (A2) defines the evidential direction at issue: discrimination consists in independently supported reasons for expecting the record more under H than under N . It does not assert that every such reason is decisive overall.

Central inference

$$C(H, N, O) \rightarrow D(H, N, O) \quad (\text{A3})$$

$$U(H, O) \rightarrow \neg D(H, N, O) \quad (\text{A4})$$

$$U(H, O) \quad (\text{A5})$$

$$\therefore \neg D(H, N, O) \quad (\text{A6})$$

$$\therefore \neg C(H, N, O) \quad (\text{A7})$$

Annotation. Line (A3) makes independently supplied discrimination necessary for public confirmation. Line (A4) says unrestricted post-hoc reassignment supplies no such discrimination; it does not say that H can never receive independent constraints. Line (A5) is the case-specific diagnostic premise. Line (A6) follows from (A4) and (A5) by modus ponens. Line (A7) follows from (A3) and (A6) by modus tollens and concludes only that this evidential route has not established public confirmation.

Validity, support, and scope

Lines (A3)–(A7) are valid; soundness depends on the premises. Line (A3) states comparative confirmation. Line (A4) applies only to post-hoc auxiliaries lacking likelihood consequences. Line (A5) is defeasible by antecedent rules, a complete denominator, comparative expectations, and downward revision. The result is $\neg C(H, N, O)$, not $\neg H$; a constrained H or pastoral use falls outside $U(H, O)$.

References

- Boudry, Maarten and Johan Braeckman (2011). “Immunizing Strategies and Epistemic Defense Mechanisms.” In: *Philosophia* 39.1, pp. 145–161. DOI: [10.1007/s11406-010-9254-9](https://doi.org/10.1007/s11406-010-9254-9).
- Dawes, Gregory W. (2009). *Theism and Explanation*. London: Routledge.
- Draper, Paul (1989). “Pain and Pleasure: An Evidential Problem for Theists.” In: *Noûs* 23.3, pp. 331–350. DOI: [10.2307/2215486](https://doi.org/10.2307/2215486).
- Hempel, Carl G. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.
- Howson, Colin and Peter Urbach (2006). *Scientific Reasoning: The Bayesian Approach*. 3rd ed. Chicago: Open Court.
- Keeley, Brian L. (1999). “Of Conspiracy Theories.” In: *The Journal of Philosophy* 96.3, pp. 109–126. DOI: [10.2307/2564659](https://doi.org/10.2307/2564659).
- Lakatos, Imre (1970). “Falsification and the Methodology of Scientific Research Programmes.” In: *Criticism and the Growth of Knowledge*. Ed. by Imre Lakatos and Alan Musgrave. Cambridge: Cambridge University Press, pp. 91–196.
- Lipton, Peter (2004). *Inference to the Best Explanation*. 2nd ed. London: Routledge.
- Popper, Karl R. (1959). *The Logic of Scientific Discovery*. London: Hutchinson.
- Swinburne, Richard (2004). *The Existence of God*. 2nd ed. Oxford: Oxford University Press. DOI: [10.1093/acprof:oso/9780199271672.001.0001](https://doi.org/10.1093/acprof:oso/9780199271672.001.0001).
- Woodward, James (2003). *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.