

Infinite Flexibility

Sean Carroll, Explanatory Elasticity, and the Predictive Cost of Theistic Accommodation

Phil Stilwell

Independent Scholar

ORCID: 0009-0009-3725-8682

Revised August 2026

Central claim

When a God-concept leaves several possible divine aims open, its expectation for an outcome is the prior-weighted average across those aims. Selecting whichever aim best fits the outcome afterward may preserve compatibility, but it cannot replace that average or manufacture confirmation without independent support. A constrained God-concept can regain evidential force by fixing aims before the evidence, excluding outcomes, and accepting revision; the argument limits elastic evidential credit rather than disproving theism.


Preview: Marginalize; do not maximize

Umbrella claim	Tempting move	Required calculation	Bounded result
A God-concept leaves several mutually exclusive aims live.	After observing the world, select the aim that makes it look most expected.	Average the aim-specific expectations using weights fixed independently of this outcome.	Best fit can restore compatibility; confirmation requires the prior mixture to outperform a live rival.

The rule applies to naturalistic and theistic hypotheses alike. The focal problem is unusually mobile divine intention, not the mere use of auxiliary assumptions or the claim that physics is the only legitimate form of explanation.

Visual preview:

The Cost of Infinite Flexibility



Read the image as a comparison of evidential methods. The left panel shows an elastic hypothesis stretching around contrary outcomes by changing its purpose claim; the wide compatibility bar means that many results remain logically possible, while the flexibility penalty marks the resulting loss of discrimination. The right panel shows the alternative: specify a risky expectation before the result, test it, and let failure reduce confidence. The image does not say that every God-concept permits literally every outcome or that naturalistic models are automatically disciplined. It shows why compatibility gained through unrestricted adjustment is weaker than confirmation earned through stable constraints.

Abstract

Sean Carroll’s challenge in the 2014 God and Cosmology debate is often summarized as a complaint that theism is not physics. Its stronger form concerns explanatory elasticity: a framework that can pair almost any observed world with some possible divine intention may retain logical compatibility while losing power to discriminate among worlds. This paper makes that challenge precise without assuming that all explanation must be scientific. It distinguishes an umbrella God-concept from narrower aim-auxiliaries and derives a simple constraint from the law of total probability. If several mutually exclusive aims remain live before an outcome, the God-concept’s likelihood for that outcome is the prior-weighted mixture across those aims. Selecting the best-fitting aim after observing the outcome substitutes a maximum for that mixture and uses the same evidence twice—first to choose the aim and then to praise its fit. I call the resulting restriction the evidential conservation rule: an auxiliary introduced by the focal evidence may restore compatibility, but it cannot by itself supply the evidential advantage it was introduced to create. A five-part constraint ledger—antecedent specification, independent support, cross-case stability, exclusion power, and revision exposure—then distinguishes disciplined auxiliaries from mobile rescue. Applications to cosmic fine-tuning, suffering, divine hiddenness, and petitionary prayer show both the force and limits of the rule. The conclusion is conditional. The argument neither disproves theism nor makes naturalism the default winner. It shows what a specified God-concept must provide to earn public confirmation from a particular world: stable aims or probability weights fixed independently of that world, a comparative likelihood advantage, and genuine exposure to evidential loss.

Keywords: Sean Carroll; explanatory elasticity; theism; God-concepts; Bayesian confirmation; auxiliary hypotheses; model selection; fine-tuning

1 Carroll’s Stronger Challenge

In his 2014 debate with William Lane Craig, Sean Carroll pressed a methodological objection to the use of theism in cosmology. The objection was not simply that current physics contains no divine variable. Nor was it the crude demand that every worthwhile explanation be a physical theory. Carroll’s sharper pressure concerned the range of worlds that theological reasoning can absorb. A beginning or an eternal cosmos, simple or intricate laws, conspicuous or hidden agency, life or vast lifeless regions can each be joined to a possible divine purpose (Carroll and Craig 2014). If the relevant purpose is chosen only after the world is known, fit comes easily.

That ease can be mistaken for strength. An account appears comprehensive because it always has something coherent to say. Yet three achievements must be separated:

1. **Compatibility:** the observation does not contradict the account.
2. **Conditional fit:** the observation would be expected if the account were joined to one particular auxiliary claim.
3. **Confirmation:** the observation favors the account over a live rival because it was more expected under the account before the observation was used to adjust it.

The first two do not automatically yield the third. A detective’s broad story can remain compatible with every clue. A narrower version of that story may fit a clue well once the detective chooses the suspect, motive, and timeline that match it. The clue confirms the broad story only if the broad story, with its admissible variants properly weighted, made that clue more expected than a rival did.

This paper develops that gap. *Theism* will name the general view that at least one deity exists. A *God-concept* will name a specified conception of a deity, including whatever attributes and intentions are relevant to an evidential argument. This distinction prevents a common slide. Theism in general

is too coarse to assign many informative expectations. A particular God-concept may be much more specific, but it must then carry the risks created by that specificity.

The proposed contribution is not another free-standing complexity penalty of the form “subtract some amount for every auxiliary.” Such scores can be useful in statistical contexts, but a theological application would invite an immediate dispute over the size and meaning of the penalty. The central result here is more basic. Probability already tells us how to assess an umbrella hypothesis whose possible aims remain open: average across them using weights that were not reverse-engineered from the focal outcome. The best-fitting aim cannot simply stand in for that average.

The evidential conservation rule

An auxiliary assumption introduced because it fits the focal evidence may restore the parent hypothesis’s compatibility with that evidence. Unless the auxiliary also has support independent of the focal evidence, it cannot supply the evidential advantage it was introduced to create. The evidence cannot first select the rescue and then be counted again as a successful prediction of the rescued hypothesis.

This rule does not single out religion. Naturalistic models can be protected by adjustable initial conditions, hidden mechanisms, flexible priors, or selective data treatment. Philosophy of science has long rejected the idea that theories face evidence without auxiliaries: Quine emphasizes the web-like character of testing, Popper emphasizes risky exposure, and Lakatos distinguishes progressive development from protective adjustment that produces no new success (Quine 1951; Popper 1959; Lakatos 1970). The issue is not whether an auxiliary exists. It is how that auxiliary is constrained and what evidential work it has independently earned.

2 From Possible Aims to a Predictive Mixture

Suppose a God-concept, call it G , is compatible with four aims concerning a four-outcome situation. One aim favors outcome 1, another favors outcome 2, and so on. Before the result is known, each aim has weight 0.25. Under its favored aim an outcome has probability 0.70; under each of the other three aims it has probability 0.10. The probabilities under each aim sum to one.

After outcome 1 occurs, an advocate can point to aim 1 and say that the result had probability 0.70. But that is not the probability assigned by the umbrella God-concept. Its probability was

$$P(E_1 | G) = 0.25(0.70) + 0.25(0.10) + 0.25(0.10) + 0.25(0.10) = 0.25.$$

The same calculation applies to every outcome. Although there is always a matching aim that gives the actual result probability 0.70, the umbrella claim gives each result probability 0.25. The after-the-fact maximum is sharp; the before-the-fact mixture is flat.

This is an application of the law of total probability, not a controversial theology-specific scoring rule. When the possible aims form a mutually exclusive and exhaustive set, the probability of an observation under the parent hypothesis is the weighted sum of its probability under each aim. Bayesian approaches make comparative likelihoods central, while likelihoodist treatments emphasize how evidence favors one hypothesis over another only through differences in expected observation (Howson and Urbach 2006; Sober 2008). The formal appendix states the rule precisely and annotates every symbol.

The word *prior* here does not mean “arbitrary” or “immune to evidence.” It means before using the particular outcome currently offered as confirmation. Historical theology, moral theory, revelation claims, prior observations, and arguments about perfect agency may all influence the weights. The

important restriction is independence: the focal outcome cannot be the sole reason for choosing the aim and then also be counted as evidence made likely by that chosen aim.

Real agents can pursue several aims at once, so the word *aim* should not be read as requiring a deity to have only one purpose. When purposes combine, each A_i can represent a complete aim-profile or decision policy—for example, the relative priority assigned to relationship, creaturely freedom, regular order, and suffering prevention. Those profiles can still be divided into nonoverlapping alternatives for the calculation. The formal requirement is a partition of the live policies, not a simplistic psychology.

The relevant comparison also matters. If a rival R gives the result probability 0.05, then the mixture value 0.25 favors G over R by a likelihood ratio of five. If R gives it probability 0.50, the same evidence favors R . Nothing in the conservation rule settles those empirical numbers. It blocks only the substitution of the selected value 0.70 for the correct mixture 0.25.

Table 1: Three different evidential questions

Question	What it establishes	What it does not establish
Is E possible under G ?	Logical compatibility: G survives E .	That G expected E , or that E favors G over a rival.
Does some aim A_i make E likely?	Conditional fit of G joined to A_i .	That A_i had substantial weight before E was observed.
Is the prior mixture for E greater under G than under R ?	A comparative likelihood advantage for G over R .	A complete posterior verdict, which may also depend on prior credibility and other evidence.

This framework is deliberately modest. Probability assignments to unique cosmic outcomes can be difficult or imprecise. But imprecision does not license selecting a favorable number after the fact. If no defensible distribution over divine aims is available, the honest conclusion is that the proposed God-concept has not yet supplied a determinate likelihood. Ignorance is not a hidden likelihood advantage.

3 Auxiliary Mobility

Every serious account uses auxiliary assumptions. A telescope result depends on calibration, noise models, and observation conditions. A historical explanation depends on transmission, dating, and human motives. A personal explanation depends on beliefs, preferences, abilities, and opportunities. The presence of auxiliaries is therefore not a defect.

The key variable is *mobility*: how freely an auxiliary can be moved after an outcome in order to preserve the parent claim. A mobile auxiliary changes with the case, lacks support apart from the case it saves, excludes little, and leaves confidence in the parent claim unchanged. A disciplined auxiliary may be revised, but its revision has a cost because it was independently supported, used across cases, or tied to further expectations.

Forster and Sober connect non-ad-hocness and simplicity to predictive performance in model selection: fitted flexibility can improve agreement with known data while degrading accuracy on new data (Forster and Sober 1994). The present argument does not equate theology with curve fitting. It borrows the general lesson that accommodation and predictive success are different achievements. Dawes similarly stresses that theistic explanations should be assessed as explanations rather than excused from ordinary comparative questions merely because they invoke agency (Dawes 2009).

The following constraint ledger turns mobility into five inspectable questions.

Table 2: The auxiliary-constraint ledger

Constraint	Disciplined auxiliary	Mobile rescue
Antecedent specification	The aim or rule is stated before the focal outcome is used.	The aim is named because it matches the known result.
Independent support	Other evidence or argument supports the aim and its weight.	The saved result is the only reason to accept the aim.
Cross-case stability	Comparable cases use the same assignment rule.	Opposite cases receive different purpose labels as needed.
Exclusion power	Some feasible outcomes would be less expected or count against the account.	Nearly every result receives a favorable reinterpretation.
Revision exposure	Failed expectations lower confidence or change the model in a recorded way.	Confidence in the parent claim remains fixed while auxiliaries move.

No single failed box proves that an auxiliary is false. A new assumption may be true even if invented in response to a surprise. Scientific progress sometimes begins with an unexpected repair. What matters is subsequent discipline: the repair should generate independent consequences, travel to new cases, outperform rivals, or otherwise earn support beyond the anomaly it was designed to absorb. Lakatos’s contrast between progressive and degenerating problem shifts captures this temporal dimension (Lakatos 1970).

The ledger also explains why “unfalsifiable” can be too blunt. A probabilistic claim need not forbid an outcome absolutely in order to be testable. It can merely make that outcome much less expected. Conversely, a claim may technically exclude something remote while remaining so flexible across the live evidence that the exclusion does little work. Predictive discipline comes in degrees.

4 Why Divine Intention Is Especially Elastic

Personal explanation is not defective merely because agents can have many motives. A familiar person’s stable preferences can make an action more expected. If a friend reliably protects animals, finding that friend at an injured animal’s side may be unsurprising for a reason. Agency becomes evidentially useful when character, commitments, circumstances, and available actions constrain the space of intentions.

A highly general God-concept often removes those constraints. Unlimited power expands the action set. Perfect knowledge removes informational limits. Hidden reasons expand the motive set. Appeals to goods beyond human understanding weaken observers’ ability to rank aims. If every apparent tension can be answered with some unknown purpose, the model’s action distribution becomes hard to specify.

This produces a dilemma rather than a simple refutation.

1. If divine aims are specified strongly enough to make some observations much more expected than others, the God-concept becomes evidentially informative. It can gain confirmation from success, but it must also risk evidential loss from failure.
2. If divine aims remain open enough to fit whichever world appears, the God-concept is protected from many failures. It also loses permission to claim that the actual world was especially expected.

Swinburne offers a major alternative to unlimited flexibility. His cumulative case assigns substantive attributes and prior simplicity to a theistic hypothesis, then asks how various observations affect its probability (Swinburne 2004). The present argument does not assume that this project fails. It identifies a burden within it. Simplicity of the agent description does not by itself specify a probability

distribution over the agent's possible choices. "A perfectly good and free deity" may be a linguistically compact description while still leaving disagreement about which worlds such an agent would create.

Nor does inference to the best explanation remove the burden. Explanatory comparison legitimately considers unification, mechanism, simplicity, and understanding, not likelihood alone (Lipton 2004). Yet an account cannot count its ability to reinterpret every result as both breadth and risky success. If fit was achieved through a mobile auxiliary, that fact remains relevant to whether the explanation is better.

5 Four Applications

The conservation rule is clearest when applied locally. It does not assign final probabilities. It asks which intention claim is doing the work, when that claim was fixed, and what other outcomes it would have made less expected.

5.1 Cosmic fine-tuning

Suppose a God-concept includes a strong antecedent aim to create embodied conscious life. A life-permitting range of constants may then be more expected under that concept than under a rival that distributes probability broadly across possible constants. This can be a genuine likelihood argument.

Elasticity enters when the aim is widened after objections. If no life had appeared, perhaps a deity valued elegant empty laws, noncarbon minds, inaccessible regions, or creation for reasons unrelated to observers. These possibilities may be coherent, but leaving all of them live changes the parent likelihood. The probability of life-permitting conditions must be averaged across the aims and their prior weights. One may not use a life-centered aim when life is observed while reserving non-life aims to protect the same God-concept in sterile worlds.

The result is not that fine-tuning provides no evidence. It is that the argument's force depends on a specified God-concept, a defensible distribution over divine choices, an observation-selection analysis, and a live rival. General theism inherits no automatic high likelihood merely from containing a purpose that matches our world.

5.2 Suffering and flourishing

The observed distribution of pain and pleasure has been used in likelihood-based evidential arguments, notably Draper's comparison of theism with a hypothesis of indifference (Draper 1989). Theistic replies cite free agency, soul-making, reliable natural order, justice, relationship, or goods beyond human grasp. Any of these may be part of a serious model. The ledger asks whether their application is stable and independently constrained.

If recovery is attributed to mercy while an otherwise similar death is attributed to a higher plan, both cases may remain compatible with the same worldview. But the paired interpretations do not automatically show that either outcome confirms it. The prior mixture must state how strongly the God-concept expected recovery, death, or their observable distribution.

Skeptical theism sharpens the issue by challenging human ability to infer that no sufficient divine reason exists merely because none is visible (Bergmann 2001). That challenge can limit confident negative likelihood judgments. It has a symmetrical consequence: if human reason is too limited to rank divine purposes in adverse cases, the same limitation restricts confident positive claims that favorable cases were antecedently likely because of a divine purpose. Mystery may block a disconfirmation argument, but it cannot simultaneously function as a sharply directed confirmation model.

5.3 Divine hiddenness

A relationship-centered God-concept might predict widespread, noncoercive awareness of the deity. Observed nonbelief then creates pressure. An advocate may add aims concerning freedom, soul-making, epistemic distance, cultural development, or the value of searching. Again, the question is not whether these are possible. It is whether they were supported and weighted before hiddenness was used to select among them.

An aim of unmistakable relationship and an aim of substantial epistemic distance pull predictions in different directions. If both remain available without weights, the umbrella concept can fit clarity and ambiguity. It cannot then claim the full predictive success of whichever aim matches our world. A richer theological theory could resolve the issue by specifying how the aims trade off and what degree or distribution of hiddenness follows.

5.4 Petitionary prayer and selective cases

Individual prayer narratives often have pastoral and personal meaning. Public confirmation is a separate claim. If improvement is answered prayer, no improvement is patient timing, and deterioration is an unknown good, then the interpretive system covers all outcomes. To gain public evidential force, an advocate needs a stable assignment rule, a denominator that includes comparable successes and failures, and an expectation that differs from a rival.

This requirement does not demand that agency behave like a vending machine. Human agents also answer requests selectively. But evidence for a person's responsiveness comes from patterns: known preferences, stated commitments, repeated choices, and contrasts with what happens in the person's absence. A divine-agent explanation may appeal to analogous constraints. What it cannot do is count only the selected outcome and treat unrecorded alternatives as irrelevant.

Table 3: Where elasticity enters the applications

Case	Potentially risky aim	Elastic fallback	Needed discipline
Fine-tuning	Embodied conscious life is strongly intended.	Empty elegance or unknown forms would also fulfill a purpose.	Prior weights over aims and constants; live comparator.
Suffering	Flourishing and protection are strongly intended.	Any loss serves freedom, formation, justice, or unknown goods.	Stable case rule, full distribution, revision on failed expectations.
Hiddenness	Relationship makes awareness likely.	Distance preserves freedom, search, or development.	A stated tradeoff predicting a degree and distribution of awareness.
Prayer	Some requests receive responsive intervention.	Yes, no, and wait each count as fitting responses.	Antecedent outcome rule, complete denominator, rival comparison.

6 The Strongest Replies

6.1 “A free person need not have statistical laws”

True. The argument requires neither strict laws nor exact numbers in every case. It requires only comparative constraint. We routinely learn about agents from tendencies, commitments, reasons, and counterfactual choices. “Very likely,” “more likely than under the rival,” and “this result would count

against the proposal” can be informative without false precision. Total refusal to rank possible actions protects freedom at the cost of predictive evidence from the action selected.

6.2 “The aim is supported by theology, not invented after the event”

Then it may legitimately receive prior weight. A longstanding doctrine, independent moral argument, or separate revelation claim can constrain a God-concept. The conservation rule is not hostile to such sources. It asks whether they genuinely select and weight the aim without relying on the same focal evidence. The resulting model should also be used consistently in difficult cases. Antecedence alone is not enough, but it removes the simplest post-hoc defect.

This answer can make theology riskier. If the independent source says a deity strongly seeks a recognizable outcome, contrary observations bear on either the source, its interpretation, or the God-concept. That exposure is not an embarrassment added by the present argument. It is the price of making the source evidentially informative.

6.3 “Unknown reasons defeat the probability comparison”

Unknown reasons may show that observers cannot confidently assign a low probability to an apparent evil or hiddenness. Yet a fully unspecified reason variable also prevents a confident high probability for an apparent blessing or fine-tuned feature. The limitation is symmetrical. A person may reasonably retain commitment for non-evidential, private, or cumulative reasons while declining to treat the selected event as new public confirmation.

6.4 “Naturalism is at least as flexible”

Sometimes it is. A naturalistic umbrella that permits any law, initial state, or multiverse distribution without prior weighting has the same problem. A detailed physical model gains credit only through stable parameters, successful novel consequences, error control, and comparison with alternatives. Constructive empiricism also warns against moving too quickly from empirical adequacy to truth (Fraassen 1980). The present result does not crown naturalism. It removes a special exemption for theistic accommodation.

6.5 “Unique cosmic facts cannot be predicted”

Prediction here means an evidential expectation specified independently of the record being assessed. It can concern traces of a past event, structural relations, ranges of constants, distributions within the one observed cosmos, or counterfactual differences among models. Unique events make probability assignment harder, not meaningless. Where no justified comparison is possible, suspension is preferable to retroactive certainty.

7 A Discipline Test for God-Concepts

Carroll’s challenge becomes constructive when converted into a test. An advocate who offers a feature of the world as evidence for a God-concept should answer six questions:

1. **Which God-concept?** State the attributes and aims relevant to this evidence. General theism is not yet a predictive model.
2. **Which alternatives within it?** List the live, mutually incompatible aims or action policies that would lead to different expectations.
3. **What fixed their weights?** Identify support independent of the focal observation. If only ranges can be defended, report ranges.

4. **What does the mixture predict?** Average across the live aims; do not report only the best-fitting conditional value.
5. **Compared with what rival?** State the alternative and its expectation for the same evidence at the same level of detail.
6. **What would reduce confidence?** Name feasible results that would lower the likelihood, revise the aim weights, or weaken the parent claim.

These questions do not guarantee consensus. Priors, rival choice, observation selection, and the content of perfect goodness remain disputed. The test provides something more modest and useful: it prevents the structure of the disagreement from disappearing behind the phrase “a deity could have wanted this.”

The test also clarifies when a newly added auxiliary earns rehabilitation. Suppose a surprising observation leads an advocate to propose aim *A*. At that moment *A* may be an accommodation. It becomes evidentially productive if it later predicts another result, connects independently supported commitments, survives comparable cases, or rules out alternatives that could have occurred. Accommodation is not a permanent stigma. It is an unpaid evidential debt that later performance can repay.

Model-selection theory gives this distinction a familiar shape. Flexibility often improves fit to the data used for selection, while successful constraint improves performance beyond those data (Forster and Sober 1994). Theology need not copy a statistical criterion mechanically. It should nevertheless distinguish fitting the known world from having had reason to expect that world.

The practical dilemma

A God-concept may remain broad enough to accommodate many possible worlds, or become specific enough to make the actual world evidentially distinctive. It can combine breadth and confirmation only by supplying independently supported weights that show why this world, rather than merely some possible world, was comparatively expected.

8 Conclusion

Sean Carroll’s critique is strongest when understood as a demand for evidential accounting. The problem is not that theism says too much, invokes agency, or fails to resemble a physical mechanism. The problem arises when a broad God-concept receives the predictive credit of a narrow aim selected only after the world is known.

The law of total probability supplies the basic correction. When several aims remain live, the parent likelihood is their prior-weighted mixture. The maximum conditional fit of the aim that matches the outcome is not that mixture. Replacing the mixture with the maximum turns universal availability into apparent predictive success and counts the same evidence twice.

The evidential conservation rule therefore places a bounded limit on accommodation. A purpose claim introduced by an outcome may show how the outcome could fit the worldview. It cannot, without independent support, turn that restored fit into evidence for the worldview. The auxiliary-constraint ledger makes the distinction operational: ask about antecedence, support, stability, exclusion, and revision.

None of this proves that no deity exists. It does not show that theological narratives lack personal, moral, or pastoral meaning. It does not establish that a naturalistic rival is coherent, simple, or probable. A well-specified God-concept may make successful, risky expectations and earn confirmation. But it

must accept the same bargain that gives prediction its force: the possibility of being wrong, or at least less favored, when the world differs from what its independently constrained aims led us to expect.

A Formal Reconstruction of the Central Thesis

This appendix presents the same argument twice. The first version uses ordinary-language premises. The second introduces probability symbols. Each symbolic line is followed by an annotation explaining what it says, why it matters, and what it does not assume.

A.1 Ordinary-language reconstruction

- P1.** When a parent hypothesis leaves several mutually exclusive and jointly exhaustive auxiliary aims live, its expectation for an outcome is the weighted average of the expectations under those aims.
- P2.** Evidence favors that parent hypothesis over a rival through the parent expectation, not merely through the expectation under whichever aim fits best after the outcome is known.
- P3.** Selecting a matching aim because of the outcome and then using the match as evidence of successful expectation uses the outcome both to choose the model variant and to score that chosen variant.
- P4.** Without evidence independent of the outcome that supported the selected aim or its weight, this selection does not change what the parent hypothesis predicted before the outcome.
- P5.** Therefore, a post-outcome aim can restore compatibility but cannot by itself create confirmation for the parent hypothesis.
- P6.** A God-concept with mobile aims can regain confirmation by fixing or independently weighting those aims, outperforming a live rival under the resulting mixture, and accepting evidential loss when its constrained expectations fail.

Premises P1 and P2 state the probability framework. P3 identifies the selection error. P4 states the independence condition. P5 follows from P1–P4. P6 is constructive: it explains how the criticized reasoning can be repaired rather than treating flexibility as an incurable property of theism.

A.2 Symbols

G: The umbrella God-concept being evaluated. It is more specific than theism in general but may still leave several divine aims open.

R: A live rival hypothesis evaluated against the same evidence and background information. The notation does not assume that *R* is naturalistic.

B: Shared background information accepted for the comparison.

I: Information that supports aim weights independently of the focal outcome. Theology, earlier evidence, or a moral argument may belong here if it is not inferred from that same outcome.

P: The probability function. An expression such as $P(E_k | G)$ reads “the probability of outcome E_k assuming God-concept G .”

n: The number of live aim-profiles included in the tested model.

Ω_A : The complete set of aim-profiles represented by that model.

A_i : The *i*th possible divine aim, complete aim-profile, or action policy within *G*. The set $A_1, \dots, A_n$ is assumed mutually exclusive and jointly exhaustive for the model being tested.

E_k : The *k*th possible observable outcome. The actually observed outcome is one member of this set.

w_i : The prior weight $P(A_i | G, B, I)$ assigned to aim A_i before using E_k . All weights are nonnegative and sum to one.

q_{ik} : The aim-specific likelihood $P(E_k | G, A_i, B, I)$: how expected outcome E_k is if G and aim A_i are true.

p_k : The parent likelihood $P(E_k | G, B, I)$ obtained by averaging the q_{ik} values with weights w_i .

r_k : The rival likelihood $P(E_k | R, B, I)$.

M_k : The largest aim-specific fit for E_k , defined as $\max_i q_{ik}$. This is the value found by choosing the best-fitting aim after seeing E_k .

Λ_k : The likelihood ratio p_k/r_k , which records how the observation discriminates between G and R at this stage of the analysis.

A.3 Symbolic derivation with annotations

Line 1: partition and weights.

$$A_i \cap A_j = \emptyset \ (i \neq j), \quad \bigcup_{i=1}^n A_i = \Omega_A, \quad w_i = P(A_i | G, B, I), \quad \sum_{i=1}^n w_i = 1.$$

Plain-English annotation. The possible aims are being treated as separate options within the tested God-concept, and together they cover the aim-options currently left open. The weights record how seriously each option was taken before the focal outcome was used. This line does not claim that real divine psychology must be simple or that researchers can always know exact weights. It states what a probabilistic version of the umbrella claim must represent.

Line 2: aim-specific expectations.

$$q_{ik} = P(E_k | G, A_i, B, I).$$

Plain-English annotation. Each q_{ik} asks a narrow question: if this God-concept and this particular aim were true, how expected would this outcome be? These narrow values are legitimate. The later error occurs only when the largest one is reported as though it had been the expectation of the whole umbrella claim.

Line 3: the required parent likelihood.

$$p_k = P(E_k | G, B, I) = \sum_{i=1}^n w_i q_{ik}. \quad (1)$$

Plain-English annotation. Equation (1) is the law of total probability. It says to multiply each aim's expectation by that aim's prior weight and add the results. In everyday terms: average across the live aims rather than pretending that only the successful one existed. This is the central formal constraint.

Line 4: comparative confirmation.

$$r_k = P(E_k | R, B, I), \quad \Lambda_k = \frac{p_k}{r_k}. \quad (2)$$

Plain-English annotation. The same outcome must be assessed under a live rival. If $\Lambda_k > 1$, this observation favors G over R in likelihood; if $\Lambda_k < 1$, it favors R ; if it equals 1, it does not discriminate between them. The ratio does not by itself settle final belief, because prior credibility and other evidence may still matter.

Line 5: post-outcome best fit.

$$M_k = \max_{1 \leq i \leq n} q_{ik}. \quad (3)$$

Plain-English annotation. M_k is obtained by looking across all the possible aims after outcome E_k is known and choosing the aim that fits it best. This value answers “Which available aim best accommodates what happened?” It does not answer “How strongly did the umbrella God-concept expect what happened?”

Line 6: why maximizing inflates fit.

$$p_k = \sum_{i=1}^n w_i q_{ik} \leq M_k \sum_{i=1}^n w_i = M_k. \quad (4)$$

Plain-English annotation. A weighted average cannot be greater than the largest value being averaged. Therefore the best-fitting aim will always look at least as successful as the parent model and often much more successful. Substituting M_k for p_k systematically rewards flexibility: the more aims available for selection, the easier it is to find an impressive match.

Line 7: compatibility does not entail confirmation.

$$(\forall k \exists i q_{ik} > 0) \not\Rightarrow (\forall k \Lambda_k > 1). \quad (5)$$

Plain-English annotation. The statement on the left says that every possible outcome fits at least one available aim. The statement on the right says that every outcome favors G over the rival. The crossed implication says the first does not logically guarantee the second. Universal accommodation can coexist with a flat mixture that predicts no outcome better than the rival.

Line 8: a concrete four-outcome example.

$$w_1 = w_2 = w_3 = w_4 = 0.25, \quad (q_{1k}, q_{2k}, q_{3k}, q_{4k}) = (0.70, 0.10, 0.10, 0.10),$$

$$p_k = 0.25, \quad M_k = 0.70. \quad (6)$$

Plain-English annotation. For the observed outcome, one aim gives a strong-looking fit of 0.70 while the other three give 0.10. Because all four aims were equally live beforehand, the correct parent expectation is 0.25. Reporting 0.70 would erase the three aims that did not match. If each possible outcome has its own favored aim, this same maneuver can make whichever outcome occurs look predicted.

Line 9: the role of independent information.

$$w_i = P(A_i \mid G, B, I), \quad I \not\equiv E_k. \quad (7)$$

Plain-English annotation. The symbol I marks the possible repair. Independent theology, earlier observations, or arguments may give one aim much greater prior weight than the others. The notation $I \not\equiv E_k$ means that this support must not merely be the focal outcome renamed. If independent information makes the matching aim dominant before the test, the mixture can legitimately approach that aim’s narrow likelihood.

Line 10: bounded conclusion.

$$\text{Post-hoc } A_i + E_k \Rightarrow \text{compatibility restored}, \quad \text{confirmation of } G \Rightarrow \Lambda_k > 1 \text{ using } p_k, \quad (8)$$

$\therefore M_k$ cannot replace p_k without independently justified aim selection.

Plain-English annotation. The first arrow grants the legitimate achievement of a rescue: it can show that the observation and the God-concept can coexist. The second arrow states the additional requirement for comparative evidential support. The conclusion does not say that the selected aim is false, that G is false, or that R wins. It says only that a match chosen by the evidence cannot be counted as if it had predicted that evidence.

A.4 Validity, support, and limits

The mixture identity in Equation (1) is deductively secured once the aim set is a partition and the conditional probabilities are defined. The inequality in Equation (4) follows from nonnegative weights summing to one. Equation (5) gives the central logical point: an existential fit claim has insufficient form to entail a comparative likelihood claim.

The application to a real God-concept is not deductive all the way down. It requires judgments about which aims are genuinely live, whether they are mutually exclusive at the relevant level, what independent information supports them, how finely the outcome space should be divided, and which rival is appropriate. Those are substantive philosophical and theological questions. The formalism makes them visible; it does not answer them by notation.

The conclusion is therefore strong but bounded. When aims are mobile and selected by the focal outcome, the resulting fit supplies no automatic confirmation. When a God-concept independently constrains its aims, states a comparative expectation, and remains exposed to revision, it escapes the criticism and may earn evidential support. That is precisely why explanatory risk is a virtue: it makes success capable of teaching us something.

Bibliography

- Bergmann, Michael (2001). "Skeptical Theism and Rowe's New Evidential Argument from Evil." In: *Noûs* 35.2, pp. 278–296. DOI: 10.1111/0029-4624.00297.
- Carroll, Sean M. and William Lane Craig (Feb. 21, 2014). *God and Cosmology: The Existence of God in Light of Contemporary Cosmology*. Public debate, Greer-Heard Forum, New Orleans Baptist Theological Seminary. URL: <https://preposterousuniverse.com/blog/2014/04/15/talks-on-god-and-cosmology/> (visited on 08/22/2026).
- Dawes, Gregory W. (2009). *Theism and Explanation*. London: Routledge. DOI: 10.4324/9780203876121.
- Draper, Paul (1989). "Pain and Pleasure: An Evidential Problem for Theists." In: *Noûs* 23.3, pp. 331–350. DOI: 10.2307/2215486.
- Forster, Malcolm and Elliott Sober (1994). "How to Tell When Simpler, More Unified, or Less Ad Hoc Theories Will Provide More Accurate Predictions." In: *The British Journal for the Philosophy of Science* 45.1, pp. 1–35. DOI: 10.1093/bjps/45.1.1.
- Fraassen, Bas C. van (1980). *The Scientific Image*. Oxford: Clarendon Press.
- Howson, Colin and Peter Urbach (2006). *Scientific Reasoning: The Bayesian Approach*. 3rd ed. Chicago: Open Court.
- 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. DOI: 10.1017/CB09781139171434.009.
- Lipton, Peter (2004). *Inference to the Best Explanation*. 2nd ed. London: Routledge. DOI: 10.4324/9780203470855.
- Popper, Karl R. (1959). *The Logic of Scientific Discovery*. London: Hutchinson.
- Quine, W. V. O. (1951). "Two Dogmas of Empiricism." In: *The Philosophical Review* 60.1, pp. 20–43. DOI: 10.2307/2181906.
- Sober, Elliott (2008). *Evidence and Evolution: The Logic Behind the Science*. Cambridge: Cambridge University Press. DOI: 10.1017/CB09780511806285.
- Swinburne, Richard (2004). *The Existence of God*. 2nd ed. Oxford: Oxford University Press. DOI: 10.1093/acprof:oso/9780199271672.001.0001.