

Unknowns Without Supernatural Default — 46

Explanatory Restraint, Open-World Comparison, and the Failure of Gap Inferences

Phil Stilwell

Independent Scholar

ORCID: [0009-0009-3725-8682](https://orcid.org/0009-0009-3725-8682)

Revised August 2026

Central claim

Defeat is not donation: when present natural models fail, confidence should fall in those models, but their lost support is not automatically transferred to a supernatural rival. Compared with error, misdescription, unknown mechanisms, and other unmodeled causes, a supernatural attribution gains public evidential support only when the evidence discriminates in its favor; current-model failure alone does not establish that. This blocks default gap reasoning without disproving supernatural causation or theism.

Preview: What an explanatory gap warrants

Starting point	Valid update	Missing bridge	Rational conclusion
Current natural models do not explain the event or report.	Lower confidence in the models that were expected to succeed.	Their failure does not show that the event is more expected under a supernatural target than under the residual alternatives.	Do not donate the lost support; seek evidence that distinguishes among the accounts still live.

The target is a default inference from unexplained to supernatural, not supernatural explanation in principle and not theism as a whole.


Visual preview:

Mind the Gap: The Unknown-Natural Reserve Principle

The Failure of Gap Reasoning



A Responsible Path Forward



The broken bridge depicts an unwarranted transfer from a gap to a conclusion; the intact bridge preserves a residual reserve of error, misdescription, and unknown mechanisms. It illustrates the no-transfer principle, not a promise that every unexplained event will receive a natural explanation.

Abstract

Current-model failure does not by itself warrant supernatural attribution. Defeat is not donation: confidence should fall in models that were expected to succeed, but their lost support cannot simply be transferred to a salient rival. This paper compares a specified supernatural hypothesis with a residual, open-world alternative containing error, misdescription, unknown natural mechanisms, mixed causes, and other unmodeled possibilities. Logic exposes the missing bridge; likelihood reasoning asks whether the gap is more expected under the supernatural target; the history of unconceived alternatives cautions against treating today's menu as complete; and an evidence ledger prevents a negative result from becoming unearned positive confirmation. These strands converge on the no-transfer principle: the gap alone does not establish that the event favors the supernatural target over the residual field. Absence can become comparative evidence when the search is sufficiently complete, the missing result was strongly expected, and the target predicts the resulting pattern better than the alternatives. Independent data and defensible prior commitments can therefore change the comparison. The conclusion blocks supernatural attribution by default; it neither proves that every event has a natural cause nor disproves theism.

Keywords: argument from ignorance; supernatural explanation; unconceived alternatives; Bayesian comparison; miracles; epistemic humility

1. A Blank in the Map Is Not a Label on the Territory

An event resists explanation. A medical recovery is unusually rapid; a report describes knowledge its subject apparently could not have acquired; an instrument records an anomaly; a coincidence seems too fitted to a need to be accidental. Inquiry eliminates several familiar causes. At that point, the sentence “we do not know” can feel unstable. “Therefore supernatural” offers closure.

The closure outruns the evidence. The paper's central rule is simple: *defeat is not donation*. “Unexplained” states a relation between an event and an inquirer's present explanatory resources. It does not yet state what the event is. The map has a blank; the blank is not itself a label on the territory. Eliminating a favored model can make its most salient rival feel confirmed, but a model's loss does not identify which of the remaining alternatives should receive the evidential gain.

The paper tests one rule:

The supernatural-default rule

If an event or report remains unexplained after serious natural inquiry, its unexplained status is positive reason to attribute it to a supernatural cause.

The rule is stronger than the modest claim that supernatural causation is possible. It is also stronger than the claim that a particular event, considered with all its features, might support a supernatural explanation. The rule treats a gap left by failed natural accounts as support in its own right. That transfer is the target.

The governing thesis is contrastive and conditional. Failure of current natural models bears against those models to the extent that they predicted success. The *no-transfer principle* says that this defeat does not by itself favor a specified supernatural hypothesis over the residual alternatives. Such favoring requires evidence that is more expected under the supernatural target. Gap status alone does not establish that comparison; additional evidence might.

This is not a disguised proof of naturalism. The residual alternative is not the proposition that every event has a natural cause. It is a bookkeeping category for live possibilities that the present inquiry has not responsibly excluded. Keeping that category open prevents ignorance from being credited to whichever named explanation happens to be most salient.

2. What Exactly Is Being Compared?

2.1 From a label to a hypothesis

Explanations are contrastive. To explain why one outcome occurred is partly to say why it occurred rather than relevant alternatives. Inference to the best explanation therefore requires an actual comparison among candidates, not merely dissatisfaction with one candidate (Lipton, 2004). In philosophy of religion, even an advocate of theistic explanation must ask whether a proposed divine action has identifiable explanatory virtues and how it compares with rivals (Dawes, 2009).

The word *supernatural* is too broad to do this work by itself. A usable hypothesis must specify enough to constrain expectations: the kind of agent or cause proposed, its powers and intentions, why it would produce this event, and what observations would have been less expected if it had not acted. Without such commitments, the label can be fitted after the fact to almost any outcome.

Call a sufficiently specified supernatural hypothesis the *target hypothesis*. Its relevant competitor is not merely the best-known natural model. It is the remaining live field after that model fails. This field can include:

- measurement, memory, testimony, selection, or reporting error;
- an event that has been described incompletely or inaccurately;
- a natural mechanism not yet understood or not yet applied to the case;
- interactions among familiar causes that were omitted from the model;
- deliberate manipulation or fraud where human testimony is involved; and
- other possibilities not yet conceived well enough to name.

This residual field is heterogeneous. It is not a single mechanism and should not be assigned a casually precise probability. Its function is more modest: it keeps the comparison from becoming artificially closed.

Defeat can redistribute confidence, but only across the whole live field. In a genuinely exhaustive two-way comparison, evidence against one side can favor the other. In an open-world comparison, however, failure of some named models does not collapse all remaining support onto one named hypothesis. Salience is not a likelihood advantage.

2.2 The open-world condition

A *closed-world* comparison assumes that the listed hypotheses exhaust the serious possibilities. That assumption is sometimes warranted. A controlled device may have only a few known states; a properly designed test may make its outcomes nearly exhaustive. But difficult historical, medical, cosmological, and testimonial questions often have an *open-world* structure: the present list of models may be incomplete.

The history of science supplies a general reason for caution. Stanford's problem of unconceived alternatives is not that current theories are probably false merely because future thinkers may be cleverer. It is that past inquirers repeatedly compared theories without conceiving later serious competitors (Stanford, 2006). The lesson is a defeasible restraint on claims of exhaustive comparison.

That restraint does not guarantee a future natural explanation. It does something narrower: it denies that a present failure has eliminated what has not been represented or tested. An unknown alternative is not positive evidence for naturalism. It is a reason not to treat the current hypothesis list as complete without an argument.

The reserve principle

When the serious hypothesis space is open, a responsible comparison retains a residual category for error, misdescription, unknown mechanisms, mixed causes, and other unmodeled possibilities. Defeat narrows the field; it does not identify the winner. The category receives no automatic victory, but it prevents automatic victory by a named rival.

3. A Cumulative Case Against Default Attribution

No single consideration must carry the whole argument. Four independent pressures converge on the no-transfer principle: the inference adds causal content, the gap lacks a demonstrated likelihood advantage, the hypothesis menu may be incomplete, and a vague supernatural label can be protected from every outcome.

3.1 The logical gap

From “the natural models examined so far do not explain this event,” it follows that those models have not supplied the explanation. It does not follow that a supernatural hypothesis has supplied it. The conclusion adds causal content absent from the premise.

Arguments from ignorance are not uniformly fallacious. Walton shows that their strength depends on the search procedure, the expertise and access of the searcher, the expected discoverability of the missing item, and the completeness of the relevant knowledge base (Walton, 1996). The mistake here is therefore not the use of absence as such. It is treating an open and heterogeneous remainder as though it had been searched to exhaustion.

The warranted update is consequently local: reduce confidence in the accounts that should have succeeded. Moving from their defeat to a supernatural cause requires a further bridge.

3.2 The likelihood gap

Evidence is comparative. On a likelihood approach, an observation favors one hypothesis over another when it is more expected under the first (Sober, 2008). If current-model failure would be unsurprising both when a specified supernatural cause acted and when an unknown mechanism, an error, or a misdescription occurred, then failure alone has little power to distinguish them.

This point does not require exact numbers. The claimant can ask qualitative questions. Would the target hypothesis have made this particular timing, content, distribution, and repeatability substantially more expected? Would the residual field have made those features less expected? Does the account predict riskily, or merely accommodate the observation after it arrives? Unless the answers separate the candidates, the gap remains a shared consequence rather than a discriminator. This is the no-transfer principle in likelihood form: one model’s failure is not another model’s evidence until that same evidence is shown to be more expected under the rival.

3.3 The historical constraint

Peirce described inquiry as the attempt to settle doubt by methods answerable to an external reality rather than by tenacity or preference (Peirce, 1877). Scientific history reinforces the method without promising that inquiry always succeeds. Phenomena once interpreted through limited conceptual resources have later received explanations unavailable to earlier investigators. This pattern makes explanatory humility rational.

The history cannot prove that the next anomaly will be natural. To use it that way would reproduce the overreach under criticism. It instead changes what a present failure warrants: the failure is compatible with limitations in models, instruments, background knowledge, and classification. History supplies a reason to

preserve the residual field, not evidence that the residual field must win. A supernatural conclusion therefore needs evidence beyond the recurring fact that inquiry is incomplete.

3.4 The every-outcome problem

A hypothesis earns explanatory credit partly by ruling out or disfavoring outcomes. Suppose clarity is taken as supernatural guidance, confusion as evidence of mystery, recovery as intervention, non-recovery as inscrutable purpose, confirmation as a sign, and disconfirmation as a test of faith. If every possible result can be assimilated with equal ease, the account explains by redescription rather than discrimination.

This is especially serious when the hypothesis has not specified an agent's aims independently of the result. Once the outcome is known, an intention can be invented to fit it. A flexible story may be meaningful, but flexibility reduces evidential force. The more outcomes a claim accommodates, the less any one outcome confirms it unless the accommodating assumptions have independent support. A target that cannot lose cannot claim a rival's loss as its gain.

3.5 An evidence ledger

Table 1 makes the central distinction visible: a datum can defeat one account without confirming another, and positive support appears only when the target earns a comparative advantage.

Table 1: What common gap evidence does and does not establish

Datum	What it bears against	What remains live	Warranted conclusion
A familiar mechanism does not fit.	That mechanism, to the degree it predicted a fit.	Other mechanisms, error, misdescription, mixed causes, target hypothesis.	Reject or revise the failed mechanism; do not transfer its lost credit automatically.
No current model explains the case.	The adequacy or completeness of the current model set.	Unknown mechanisms and other unmodeled causes, as well as the target.	Mark an explanatory gap and widen inquiry.
Several ordinary explanations were seriously tested and failed.	Those tested explanations.	Untested alternatives and the target; their relative weights depend on the tests.	Narrow the field without declaring it closed.
Additional features are independently predicted by the target and unlikely across the residual field.	Rival accounts that do not expect those features.	The target plus any alternatives that also predict them.	Credit the target in proportion to the discrimination and the reliability of the data.

The last row matters. The no-transfer principle is not a rule that a supernatural hypothesis can never gain evidence. It identifies what such evidence must do: distinguish the target from the whole field still live.

4. When Absence Can Become Comparative Evidence

The no-transfer principle is a constraint on automatic inference, not a rule that negative evidence can never support a rival. A missing result becomes informative when it was strongly expected and the search was capable of finding it. If a smoke alarm is working, sensitive, correctly placed, and expected to sound in the presence of dense smoke, its silence is evidence against dense smoke. If a complete database would contain a record were a transaction genuine, a careful search that finds none is evidence against genuineness. The force comes from the expected detectability of the missing item and the completeness of the search.

The same structure appears in debates over miracles. Hume's discussion links the credibility of testimony to competing explanations of the report and to the evidential background (Hume, 1748/2000). Earman argues that no blanket maxim can settle every miracle claim and that the evidence must be treated probabilistically rather than dismissed in advance (Earman, 2000). This is a useful constraint on both sides. One may not declare supernatural causation impossible by definition; one may not declare it established because familiar causes have run out.

For a gap to become comparative support for a target, at least four conditions should be addressed:

1. **Expected success.** The rejected account would probably have explained or detected the result if it were true.
2. **Search quality.** The investigators had suitable methods, access, controls, and competence.
3. **Coverage.** The alternatives treated as excluded were actually represented by the search; an open remainder was not silently erased.
4. **Target discrimination.** The observed features are more expected under the specified target than under the alternatives left live.

Meeting the first three can make the absence evidential against a class of alternatives. The fourth is needed before that negative evidence becomes comparative support for the target. Without it, the result may justify "not these accounts" while leaving "which account?" unresolved. The exception therefore confirms the central theme: evidential transfer is earned by comparative discrimination, not created by an unexplained remainder.

5. How a Supernatural Hypothesis Could Earn Support

The no-transfer principle is not a veto on supernatural explanation. It replaces an automatic inference with a burden of comparison. A serious supernatural explanation can meet that burden, but it should be evaluated as an explanation rather than protected as an exception to explanatory standards. Five tasks are central.

5.1 Specify the target before fitting the result

State what kind of cause is proposed, what capacities and intentions it has, and why those features are independently warranted. A God-concept is a particular conception of a deity, including attributed powers, intentions, character, and relations to the world; *theism* is the broader claim that at least one deity exists. Neither term by itself predicts this event at this time. The more specific attribution needs its own bridge.

5.2 Compare with a genuine residual alternative

Do not let one named natural model stand for every non-supernatural possibility. Partition the field as well as the evidence permits, and retain a residual category when it does not. If no defensible probability can be assigned to that category, report a range of results or a qualitative sensitivity judgment rather than manufacturing precision.

5.3 Use independently secured data

Testimony and anomaly reports require controls for transmission, memory, selection, prior publicity, multiple access to the same information, and incentives. Independence matters because ten reports copied from one source are not ten independent observations. Extraordinary content does not make a source unreliable by definition, but it raises the importance of tracing how the evidence was produced.

5.4 Seek discriminating prediction or retrodiction

A retrodiction explains evidence already known; it can still be informative if the account was constrained independently and would not fit many contrary results. Stronger evidence arises when a hypothesis identifies

a pattern before it is inspected or exposes itself to a result that would count against it. The question is always contrastive: what would be substantially less expected if the target were false?

5.5 Show sensitivity to priors and the residual field

Swinburne develops a cumulative probabilistic case for theism in which prior considerations and many strands of evidence jointly affect the result (Swinburne, 2004). Whether that case succeeds is beyond the present argument, but its form illustrates an important point: a conclusion cannot be read from a gap alone. It depends on how the target is specified, how likely the evidence is under rivals, and what prior commitments are defended.

A responsible presentation should therefore say which assumptions drive the conclusion. If modest changes to the prior odds or the weight assigned to the residual field reverse the result, confidence should reflect that sensitivity. If the result survives a broad and defensible range, the case is stronger.

The no-transfer test

Name the target. Preserve the residual field. Trace the data. Ask what the target predicted that the alternatives did not. Vary the defensible assumptions. State what would defeat the conclusion.

6. Objections and Replies

6.1 “The reserve principle smuggles in naturalism”

It does not assign the residual field a naturalistic guarantee. Error and misdescription are not deep metaphysical theories, and “unknown mechanism” does not mean “mechanism that must ultimately be natural.” The reserve is a refusal to treat an incomplete list as exhaustive. A well-supported supernatural hypothesis can outperform it; a gap does not do so automatically.

6.2 “If natural and supernatural exhaust the possibilities, defeating one supports the other”

That conditional is correct only when the partition is exhaustive and the evidence bears against one side as a whole. Failure of current named models establishes neither condition. It defeats the models that should have succeeded, not every natural mechanism, every error process, every misdescription, and every mixed cause. Even in a closed comparison, the size of the update depends on how strongly each side expected the evidence. A broad metaphysical dichotomy therefore does not validate a local leap from “these models failed” to “this supernatural cause acted.”

6.3 “An unknown alternative has no likelihood”

A completely unspecified remainder does not yield a neat likelihood, and pretending otherwise would be false precision. But the defect does not justify setting its weight to zero. In open-world inquiry, one can use ranges, partitions, robustness checks, and qualitative comparisons. The honest conclusion may be that the posterior comparison is underdetermined. That is a limitation of the evidence, not a victory for the named hypothesis.

6.4 “Enough extraordinary evidence can overcome low prior odds”

Yes. Bayesian comparison does not impose a permanent veto. Evidence with a sufficiently large likelihood advantage can overcome low prior odds, and independent pieces of evidence can accumulate. Earman’s criticism of overly sweeping anti-miracle arguments is pertinent here (Earman, 2000). The present claim is only that failure of current natural models does not itself establish the needed likelihood advantage.

6.5 “Every theory begins in a gap”

Gaps often begin inquiry. An anomaly can motivate a hypothesis, direct experiments, or reveal that a model needs revision. Discovery and justification are different stages, however. A gap can suggest a supernatural hypothesis without confirming it. The hypothesis earns confirmation through what happens next: specification, comparison, risky testing, and survival of defeaters.

6.6 “The supernatural may not be lawlike”

Agency explanations need not predict by exceptionless laws. Historians and juries reason about persons by integrating capacities, intentions, circumstances, and traces. But agency still imposes constraints. An account of deliberate action must say why this action, in this setting, fits the attributed agent better than alternatives. If no possible pattern counts against the attribution, calling the cause personal has not solved the every-outcome problem.

These objections change how the comparison should be conducted; they do not reverse its direction. A supernatural target may earn support through specification, reliable data, comparative fit, and defensible priors. It does not earn that support merely by inheriting the failure of accounts that were already expected to be incomplete.

7. A Reader’s Inquiry Protocol

When an event is presented as inexplicable and therefore supernatural, the following sequence applies the no-transfer principle and keeps the reasoning transparent:

1. **Describe before interpreting.** Separate the event, report, measurements, and uncertainties from the proposed cause.
2. **List what was actually tested.** Name the models, controls, and expected results rather than saying only that “science cannot explain it.”
3. **Record what each failure defeats.** Reduce confidence in the accounts that should have succeeded; do not donate the reduction to a rival.
4. **Keep an explicit residual category.** Include error, misdescription, unknown mechanisms, interactions, and unconceived possibilities as appropriate.
5. **Specify the supernatural target.** State the proposed cause’s relevant powers, intentions, and independently supported constraints.
6. **Identify discriminators.** Ask which features are more expected under the target and why those features are reliable and independent.
7. **Stress-test the comparison.** Vary plausible priors, data-quality judgments, and residual weights. Note where the conclusion changes.
8. **State the stopping rule.** Say what further result would confirm, weaken, or defeat the attribution.

This protocol does not guarantee agreement. It makes disagreement more informative. One party may defend a higher prior for theism, another may assign greater weight to testimonial error, and both may disagree about independence. Those are real points of dispute. “There is a gap” conceals them. The decisive question is: what evidence favors this target over every alternative that remains live, rather than merely disfavoring one rejected model?

8. Conclusion

An explanatory failure is information, but defeat is not donation. When a model expected success and fails, confidence in that model should fall. When several models fail, the known field narrows. Neither result alone identifies which remaining account should receive the evidential gain.

The case against supernatural default is cumulative. Logic exposes an added conclusion; likelihood reasoning demands a comparison; the history of unconceived alternatives cautions against a prematurely closed menu; and the every-outcome problem shows why an unconstrained label can accommodate without explaining. The corrective is not a promise that nature will eventually explain everything. It is open-world bookkeeping joined to discriminating tests.

A specified supernatural hypothesis may receive support from independent, reliable, and comparatively favorable evidence. Defensible prior commitments may matter, and strong cumulative evidence may overcome low prior odds. Until those elements are supplied, “unexplained” warrants continued inquiry and calibrated uncertainty. The reader should therefore ask not merely why current models failed, but what evidence favors the proposed target over the alternatives still live. A blank in the map remains a reason to investigate, not permission to write the desired destination into the terrain.

A. Formal Reconstruction of the Central Thesis

The central inference can first be stated without assigning exact numbers:

1. **Premise 1.** Public evidential warrant for a specified supernatural hypothesis over the residual live alternative requires evidence that discriminates in favor of the former relative to the latter.
2. **Premise 2.** Failure of current specified natural models is compatible with both the supernatural target and the residual alternative, and the failure alone does not establish that it is more expected under the target.
3. **Premise 3.** When the available evidence does not discriminate between the target and the residual alternative, inferring the target merely because named current models failed is a contrastive error.
4. **Conclusion.** Therefore current-model failure alone does not warrant defaulting to the supernatural target over the residual alternative; independent data and defensible prior commitments may alter that result.

A.1 Symbols

- S : a specified supernatural hypothesis.
- R : the residual live alternative, including error, misdescription, unknown natural mechanisms, mixed causes, and other unmodeled possibilities.
- G : the information that the current specified natural models fail to explain the event or report.
- D : additional data offered as discriminating evidence.
- P : a probability function representing the comparison under stated assumptions.
- B_G : the gap-only Bayes factor comparing the expectedness of the model failure under the two live sides.
- W_G : the claim that the gap information G favors S over R .

A.2 Posterior odds and the central inference

Where the relevant probabilities can be represented, posterior odds satisfy:

$$\frac{P(S | G, D)}{P(R | G, D)} = \frac{P(G, D | S)}{P(G, D | R)} \times \frac{P(S)}{P(R)}. \quad (1)$$

For the gap information considered alone, define:

$$B_G = \frac{P(G | S)}{P(G | R)}. \quad (2)$$

The gap raises the odds of S over R only insofar as B_G exceeds one. But G says that specified current models failed; it does not by itself establish that failure was more expected under S than under R . In compressed form, the no-transfer result is:

$$W_G \iff (B_G > 1), \quad G \not\Rightarrow (B_G > 1), \quad \therefore G \not\Rightarrow W_G.$$

Additional data D can change the comparison when they are independently reliable and more expected under S than under R .

Validity. The conclusion follows if the premises are granted: if comparative discrimination is required, and G does not establish it, then G alone does not warrant defaulting to S .

Premise support. The main disputes concern whether the target is specified enough to generate expectations, whether the residual field has been responsibly represented, whether the data are reliable and independent, and whether the gap itself has a likelihood advantage. The reconstruction does not prove S false, prove R true, or require exact probabilities where the evidence does not sustain them.

References

- Dawes, G. W. (2009). *Theism and explanation*. Routledge. <https://doi.org/10.4324/9780203876121>
- Earman, J. (2000). *Hume's abject failure: The argument against miracles*. Oxford University Press. <https://doi.org/10.1093/0195127382.001.0001>
- Hume, D. (2000). *An enquiry concerning human understanding* (T. L. Beauchamp, Ed.). Clarendon Press. (Original work published 1748)
- Lipton, P. (2004). *Inference to the best explanation* (2nd ed.). Routledge.
- Peirce, C. S. (1877). The fixation of belief. *Popular Science Monthly*, 12, 1–15. <https://www.peirce.org/writings/p107.html>
- Sober, E. (2008). *Evidence and evolution: The logic behind the science*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511806285>
- Stanford, P. K. (2006). *Exceeding our grasp: Science, history, and the problem of unconceived alternatives*. Oxford University Press. <https://doi.org/10.1093/0195174089.001.0001>
- Swinburne, R. (2004). *The existence of god* (2nd ed.). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199271672.001.0001>
- Walton, D. N. (1996). *Arguments from ignorance*. Pennsylvania State University Press.