

Abductive Inference and Epistemic Closure in Christian Apologetics — 21

Phil Stilwell

Independent Scholar

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

Revised August 2026

Central claim

Christian apologists who move from generic theism and a “best explanation” judgment to high confidence that a particular miracle occurred have not completed the inference. Generic theism neither entails a resurrection nor fixes its prior probability, while explanatory ranking does not by itself supply dependence-aware likelihoods across a serious rival set. Without that calibration, the miracle remains a candidate rather than a publicly warranted high-credence conclusion; strong dependence-aware evidence or a specified intervention model could change the result.

Preview:

Target move	Required bridge	What IBE supplies	Public conclusion
Generic theism plus “best explanation” is said to yield high confidence in a resurrection.	Either entailment or evidence strong enough to cross a stated probability threshold is needed.	A ranking among considered rivals; no fixed prior, rival mixture, or independence structure follows.	No warrant-transmitting closure occurs. Calibrated historical evidence could still support the claim.

The target is a public evidential inference to a specific miracle, not every form of theism, private religious warrant, or the legitimacy of abductive reasoning in well-calibrated domains.

Visual preview:

TWO BRIDGES STILL REQUIRED



Candidate now; conclusion only after calibration.

The diagram separates explanatory ranking from the two bridges required for public high credence: a specified prior under the relevant theological background and a robust Bayes factor built from serious rivals and dependence-aware evidence. It shows a burden of inference, not the impossibility of a miracle.

Abstract

Generic theism and the judgment that resurrection is the best explanation of selected historical evidence do not, by themselves, publicly warrant high confidence that a resurrection occurred. This paper separates three jobs often compressed into one apologetic move: generating a candidate, ranking it against considered rivals, and calibrating confidence. It then compares the logic of epistemic closure with Bayesian odds. The main result is that no closure principle applies: generic theism does not entail a particular miracle, and an abductive ranking neither fixes the prior odds nor establishes the Bayes factor needed to cross a high-credence threshold. Rival-set incompleteness and dependence among testimonial sources further make an uncalibrated ranking fragile. A sensitivity table replaces arbitrary miracle probabilities with explicit evidential burdens, and a reconstruction shows exactly what a stronger case would have to supply. The conclusion is bounded. Abduction remains valuable; explanatory relations can guide confirmation; low prior probability is a burden rather than a veto; and sufficiently discriminative, dependence-aware evidence could support a miracle claim. What fails is the warrant-transmitting shortcut from generic theism and explanatory appeal to public high credence in one specific intervention.

Keywords: abduction; inference to the best explanation; epistemic closure; Bayesian confirmation; miracle claims; resurrection; testimonial dependence

1. The Inference Under Pressure

Christian apologetics often presents a cumulative movement rather than a single argument. Natural theology is said to make theism credible; historical considerations are then said to make a resurrection the best explanation; confidence in the specific miracle is treated as the rational completion of the case. The movement can sound seamless because each step uses familiar epistemic language: explanation, evidence, warrant, and cumulative support. Yet the final conclusion is far more specific than the starting point.

The intended audience here is the Christian apologist who offers resurrection as public evidence—evidence meant to carry weight for a reader who does not already accept Christian revelation. The belief under challenge is not theism and not even the bare possibility of a miracle. It is the inference that generic theism plus an abductive ranking can transmit high public warrant to one particular historical intervention.

By *public high credence*, I mean a level of confidence offered as rationally warranted by shareable evidence for an inquirer who does not begin with the Christian conclusion. The phrase does not name one mandatory numerical cutoff. It does require the advocate to identify a confidence target and to show that the evidence can reach it across the important disputed assumptions. That is the calibration burden developed below.

A *God-concept* is a specified conception of a deity: not merely that some deity exists, but what capacities, intentions, policies, and relations to human history are attributed to that deity. Generic theism leaves many of those features open. A resurrection hypothesis needs more. It requires a deity able and willing to intervene, a reason to intervene at this time and place, a reason to select this person and this mode of vindication, and a reason to leave the particular evidential record under review. Those additions may be defensible. They cannot be smuggled into the word *theism* and then counted as though generic theism supplied them for free.

The warrant-transmission test

If a starting commitment is supposed to transmit public warrant to a more specific claim, ask which bridge carries the warrant. A deductive bridge requires entailment. A probabilistic bridge requires prior odds, likelihoods for the evidence under the target and its rivals, and a justified treatment of dependence. If neither bridge is supplied, “best explanation” names a preference without yet calibrating confidence.

This paper sharpens the skeptical pressure while correcting two tempting overstatements. First, adding a class of unknown alternatives to a Bayesian denominator does not mathematically guarantee that a posterior

falls; the result depends on the probability assigned to that class and on how well it predicts the evidence. Second, a low prior is not an automatic defeat. Sufficiently strong evidence can overcome low prior odds. The responsible objection is therefore not that miracles are impossible or that Bayes' theorem always favors naturalistic explanations. It is that the advertised high confidence requires an evidential bridge that the abductive label alone does not provide.

2. Three Jobs Hidden Inside “Best Explanation”

2.1 Generation, comparison, and calibration

Abductive reasoning performs at least three distinguishable jobs. It can *generate* a hypothesis worth investigating. It can *compare* candidates by explanatory scope, fit, simplicity, mechanism, or coherence. It can also be offered as a rule of *confirmation*: infer, or raise confidence in, the best explanation. Harman's classic formulation gave inference to the best explanation its modern name, while Lipton developed a sophisticated account of how explanatory considerations guide induction (Harman, 1965; Lipton, 2004). Douven emphasizes that IBE is better treated as a family of possible rules than as one fully specified algorithm (Douven, 2017).

These jobs need not succeed together. A candidate can be fruitful without being probable. One explanation can be better than the alternatives considered while all candidates remain poorly supported. A hypothesis can also explain the evidence well while having low posterior probability because the background makes it antecedently unlikely or because a broad rival class explains the evidence almost as well.

The key distinction is therefore between *best of those considered* and *sufficiently supported*. An apologetic argument that scores resurrection against hallucination, conspiracy, and apparent death may establish a local ranking if the comparisons are fair. That ranking does not yet show that the rival set is adequate, that the evidence is independent, or that the winning hypothesis crosses any stated confidence threshold.

2.2 Explanation can matter without dictating the posterior

The argument does not require the crude claim that explanation is evidentially irrelevant. Probabilistic measures can represent explanatory power, and explanatory relations can change how evidence bears on a hypothesis (Climenhaga, 2017; Schupbach & Sprenger, 2011). Lipton's own discussion seeks friendship between Bayesian and explanatory approaches rather than a simple replacement of one by the other (Lipton, 2004). The problem begins when an explanatory virtue is treated as an unpriced warrant voucher.

Suppose a hypothesis would, if true, make the evidence much less surprising. That is relevant to the likelihood $P(E | H)$. But confirmation is comparative. What matters is how much less surprising the evidence is under H than under relevant alternatives. Even a high $P(E | H)$ does not yield a high posterior when $P(E | \neg H)$ is also high or the prior odds are very low. Explanatory judgment must therefore be connected to the probability terms it is claimed to improve.

Norton's material theory supplies a further caution. Inductive success depends on facts within a domain, not merely on a universal formal virtue such as simplicity or explanatory appeal (Norton, 2021). Scientific uses of IBE often operate within networks of measurement, intervention, replication, and background theory that discipline which explanations are serious and which virtues track truth. Historical miracle claims may still be assessed rationally, but they do not inherit that calibration merely by sharing the phrase “best explanation.”

3. Why Epistemic Closure Does Not Complete the Move

Epistemic closure is usually expressed with an entailment: if a person knows p , knows that p entails q , and competently deduces q , then the person knows q . Dretske's influential challenge concerns whether knowledge is always closed under known entailment, especially when a conclusion raises relevant alternatives not settled by

the original evidence (Dretske, 1970). Whatever position one takes in that debate, closure requires an implication to close over.

Generic theism does not entail a bodily resurrection, much less the resurrection of one specified person in first-century Judea. A deity could refrain from intervention, intervene differently, reveal through another person, or have purposes that do not include public historical signs. The conditional “if theism, then this resurrection” is not a theorem of theism. It is a substantive theological proposal.

This means the apologetic move fails as closure before disputes about anti-closure even begin. The missing link is not a philosophical exception to an otherwise valid entailment; there is no entailment. At most, theism changes the background probability of the miracle. Once the claim is stated that way, the issue is Bayesian or otherwise inductive: how much does this *specified* theological background raise the prior, and how strongly does the historical evidence discriminate the miracle from its rivals?

Plantinga’s account of warranted Christian belief does not supply the missing public bridge. His proper-function epistemology is externalist: on the extended Aquinas/Calvin model, Christian belief could be warranted when produced by properly functioning faculties in an appropriate environment under a truth-aimed design plan (Plantinga, 2000). That is a serious proposal about warrant. It is not a demonstration that shared historical evidence gives a non-Christian inquirer high posterior confidence in a resurrection. A belief could be properly basic on that model while the public evidential case remains unsettled. Keeping these projects distinct prevents private warrant from being counted twice—first as background confidence and again as public confirmation.

4. The Bayesian Burden Made Visible

Let M be the proposition that a specified miracle occurred, E the public historical evidence, and T the relevant theological background. Bayes’ theorem is clearest in odds form:

$$\underbrace{\frac{P(M | E, T)}{P(\neg M | E, T)}}_{\text{posterior odds}} = \underbrace{\frac{P(E | M, T)}{P(E | \neg M, T)}}_{\text{Bayes factor}} \times \underbrace{\frac{P(M | T)}{P(\neg M | T)}}_{\text{prior odds}}. \quad (1)$$

The equation does not dictate the inputs. It disciplines the inference by showing where disagreements live (Sober, 2008; Sprenger & Hartmann, 2019). Generic theism does not fix $P(M | T)$ because the term depends on the specified deity’s intervention policy and purposes. “Best explanation” does not fix the Bayes factor because that ratio depends on the joint evidence under both the miracle and the composite alternative.

Instead of pretending to know the correct miracle prior, Table 1 displays the evidential burden under several hypothetical priors. The entries are thresholds, not estimates. For a target posterior q , the required Bayes factor is

$$B_{\min} = \frac{q}{1 - q} \times \frac{1 - p}{p}, \quad (2)$$

where $p = P(M | T)$.

Hypothetical prior p	Bayes factor needed for 50% posterior	Bayes factor needed for 90% posterior
10%	9	81
1%	99	891
0.1%	999	8,991

Table 1: Sensitivity thresholds for illustrative prior probabilities. The table does not assign an actual prior to a resurrection. It shows why a claim of high posterior confidence must defend both its theological prior and its evidential Bayes factor.

A low prior is therefore a burden, not a prohibition. Earman’s critique of simple Humean arguments stresses that sufficiently strong and sufficiently independent testimony can, in principle, overcome a low prior (Earman, 2000). Swinburne’s case correspondingly does not rely on generic theism alone; it develops natural-theological and theological reasons intended to raise the relevant background probability before the historical evidence is applied (Swinburne, 2003). The proper critical question is whether those reasons and the historical likelihoods are strong enough, not whether Bayes’ theorem excludes miracles by definition.

The same charity sharpens the objection. If the argument needs a detailed interventionist God-concept to raise the prior, that model must be stated and defended. If it needs large Bayes factors, the likelihoods must be justified against serious alternatives. Without both, high confidence is an asserted conclusion rather than the result of a completed probability comparison.

5. Rival Space Without a Mathematical Shortcut

Let $R_1, \dots, R_n$ be specified non-miracle explanations and let U collect mechanisms not represented by those rivals. Conditional on $\neg M$ and T , the alternative likelihood is a mixture:

$$P(E \mid \neg M, T) = \sum_{i=1}^n P(E \mid R_i, T)P(R_i \mid \neg M, T) + P(E \mid U, T)P(U \mid \neg M, T). \quad (3)$$

This formula corrects an easy mistake. Merely inserting U does not guarantee that $P(M \mid E, T)$ decreases. The effect depends on the mass moved into U and on whether U makes the evidence more or less expected than the rivals whose weight it displaces. Nor is “all unconceived alternatives” a neatly measurable scientific kind with an obvious likelihood.

The legitimate lesson is robustness. A posterior can be overstated when a few handpicked rivals are treated as exhaustive and their weights are renormalized to fill the entire alternative space. Stanford’s historical argument about unconceived alternatives targets precisely the confidence drawn from outperforming only the theories a community happened to conceive (Stanford, 2006). Its direct evidence concerns theoretical science, so applying it to ancient miracle inference is analogical rather than automatic. Jantzen’s analysis supplies a more specific historical caution by arguing that the testimony available for miracle assessment is affected by sampling and survival processes (Jantzen, 2012).

The practical response is not to invent a large probability for an amorphous unknown class. It is to perform a sensitivity analysis. Vary the weight and evidence-likelihood of an “other mechanisms” reserve across a defensible range; report whether the qualitative result survives. A conclusion that collapses as soon as a modest reserve is admitted is not robust enough to bear high public confidence.

The central distinction therefore survives the correction: a resurrection may rank first among named rivals while the likelihood comparison remains too fragile to calibrate public high credence. Local explanatory rank and robust comparative support are different achievements.

6. Dependence: Evidence Must Be Joined Before It Is Multiplied

Cumulative cases often gain their force by multiplication. If several items of evidence are conditionally independent, their joint likelihood can be factored:

$$P(E_1, \dots, E_m \mid M, T) = \prod_{j=1}^m P(E_j \mid M, T), \quad (4)$$

with an analogous product under $\neg M$. But the factorization is licensed only by conditional independence, and independence must be considered under both sides of the comparison. Shared texts, common oral traditions, later harmonization, community incentives, and mutual influence can make reports less than independent. Dependence under the miracle hypothesis can matter too; no general correction is obtained by adjusting only the skeptical denominator.

The McGrews' cumulative resurrection argument recognizes this issue explicitly and devotes attention to the independence of its evidence streams (McGrew & McGrew, 2009). The methodological demand is therefore not that testimony can never accumulate. It is that an advocate may not treat three dependent reports as three fresh multipliers merely because they appear in separate textual locations.

When the dependence structure is uncertain, the honest options are straightforward: model a range of correlations, combine tightly linked reports into one evidence unit, or present upper and lower bounds. Any of these is more informative than silently assuming independence. A high posterior that appears only under the strongest multiplicative assumption should be reported as fragile.

Here too the governing point is about the missing bridge. A local IBE result can remain intact even when dependence sharply limits what that result can support about the magnitude of confidence.

7. An Evidence Ledger for a Public Resurrection Case

Table 2 shows the smallest complete set of burdens. It is not a substitute for historical work; it prevents one kind of success from doing the work of another.

Burden	What must be supplied	Common shortcut	Bounded conclusion if unmet
Theological prior	A specified interventionist God-concept and reasons that make this event, person, time, and signal antecedently expected	Treating generic theism as though it entailed or strongly predicted resurrection	The prior remains open; theism alone does not close the gap
Rival space	Serious historical, psychological, textual, and social mechanisms plus robustness to an “other” reserve	Defeating a short list and treating it as exhaustive	The winning rank is local to the considered set
Joint likelihoods	Evidence probabilities under the miracle and the composite alternative	Counting explanatory appeal as a large but unstated Bayes factor	Directional support may exist, but its magnitude is uncalibrated
Dependence	A source and transmission structure under both sides of the comparison	Multiplying related reports as independent witnesses	The cumulative force is an upper bound, not an established result
Threshold	A stated level of confidence and a sensitivity analysis across disputed inputs	Calling the result “probable” or “near certain” without a burden calculation	High public credence has not been shown

Table 2: The public evidential burdens are cumulative. Meeting one does not waive the others, and failure to meet them does not prove that the miracle did not occur.

This ledger exposes a symmetry that apologetic rhetoric can obscure. Naturalistic rivals should not receive lenient priors, vague mechanisms, or double-counted evidence. Neither should the miracle hypothesis. The same standards apply: specify the model, predict the evidence, justify the comparison class, disclose dependence, and test sensitivity.

It also shows why the conclusion is less sweeping than a general rejection of IBE. In a mature domain, a candidate's mechanisms may be independently established, the rival space may be constrained by intervention and measurement, and repeated evidence may support a stable likelihood model. Under those conditions explanatory ranking can be well calibrated. In a singular ancient case, the inferential scaffolding is harder to establish. Harder is not impossible; it means the scaffolding must be shown rather than borrowed from scientific success elsewhere.

8. Strong Replies and Their Costs

8.1 “Theism radically changes the prior”

This reply can be correct in form. If a well-supported theological model gives the deity under examination a strong reason to raise Jesus bodily and to leave the observed record, then $P(M \mid T)$ may be much higher than a frequency of ordinary human revivals would suggest. Swinburne argues along exactly this route (Swinburne, 2003). The cost is specificity. The advocate must defend not only a miracle-capable deity but the intentions and historical policy that make this resurrection expected. A generic creator or ground of being does not supply those details.

That concession confirms rather than escapes the paper's architecture: a specified theological model can supply part of the probabilistic bridge, whereas generic theism cannot do so by itself.

The reply also risks circularity when the Christian story used to raise the resurrection prior is itself accepted largely because of the resurrection evidence. The prior background must have support not wholly borrowed from the evidence whose force it is being used to amplify. Otherwise the same evidence enters once to build the theological expectation and again to confirm its fulfillment.

8.2 “Historians do not assign numbers”

Historical reasoning often remains qualitative, and fabricated decimals would create false precision. Nothing here requires a canonical numerical prior. It requires directional accountability. An advocate can state which model makes an observation more expected, identify the dependence among sources, and show that the result persists across a plausible range of priors and likelihoods. Sensitivity analysis is valuable precisely when exact values are disputed.

The phrase “most probable explanation” creates a stronger burden than the phrase “a serious explanation.” If numbers cannot responsibly be assigned, the conclusion should be correspondingly modest: the miracle may explain the evidence well, may deserve consideration, or may outrank named rivals under stated judgments. It should not be advertised as near certainty on the authority of an unstated calculation.

The demand is therefore calibration in proportion to the conclusion, not numerical theater. The stronger the advertised confidence, the more clearly the bridge from explanatory rank to that confidence must be exposed.

8.3 “A cumulative case can overcome each weak strand”

Yes—when the strands are genuinely discriminative and sufficiently independent. Cumulative force is not an illusion. The ledger asks what makes it legitimate. If all strands descend from a shared narrative source, express the same community's interpretation, or are selected by the same survival process, multiplication exaggerates their combined force. If they arise independently and have low probability under serious rivals, the cumulative case strengthens. The dispute is empirical and structural, not a ban on accumulation.

8.4 “Properly basic belief does not need public evidence”

That position changes the project rather than answering it. A person may claim warranted Christian belief through a noninferential route described by Plantinga’s model (Plantinga, 2000). The present argument concerns the public use of historical evidence to move an uncommitted reader toward high confidence. A private-warrant account may coexist with the conclusion that the public abductive case is undercalibrated.

9. Conclusion

The shortcut from generic theism and explanatory preference to high confidence in a particular miracle does not work. Generic theism supplies no entailment, so epistemic closure has nothing to close over. Treated probabilistically, the move requires a specified theological prior and a Bayes factor built from serious rivals, joint likelihoods, and an honest dependence structure. “Best explanation” by itself supplies none of those magnitudes.

The cumulative case is strongest when this criticism is accepted rather than evaded. Abduction can nominate resurrection as a serious candidate. Explanatory relations can guide confirmation. A detailed theological model may raise the prior, and sufficiently discriminative, dependence-aware evidence may overcome even a substantial initial burden. But each gain must be argued in its proper place. Unknown alternatives should prompt robustness checks rather than arbitrary penalty terms; dependent testimony should be joined before it is multiplied; and high confidence should be tied to an explicit threshold.

The principal limit is therefore also the live alternative: this paper does not show that a resurrection did not occur, that theism is false, or that every evidential case for a miracle must fail. It shows that high public credence is not transmitted by the words *theism*, *abduction*, or *closure*. Christian apologists who seek public intellectual force should disclose the bridge. Readers should withhold the advertised confidence until that bridge bears the full specificity of the conclusion.

A. Formal Reconstruction of the Central Thesis

A.1 Syllogistic reconstruction

- P1.** Within a public evidential argument, high credence in a specific miracle requires either a valid entailment from warranted premises or comparative evidence strong enough, given prior odds and serious alternatives, to cross a stated high-credence threshold.
- P2.** Generic theism does not entail a particular resurrection and does not by itself fix high prior odds for that event.
- P3.** A judgment that resurrection is the best explanation among considered rivals does not by itself establish the joint likelihoods, rival mixture, or dependence structure needed to show a threshold-crossing Bayes factor.
- P4.** Therefore generic theism plus an uncalibrated abductive ranking does not publicly warrant high credence in the specific resurrection.
- P5.** This conclusion is defeasible: independently supported theological expectations or sufficiently strong, dependence-aware historical evidence could satisfy the missing probabilistic bridge.

The core inference from P1–P3 to P4 is valid because P2 removes the deductive route and P3 removes the uncalibrated probabilistic route identified in P1. P5 states the scope limit rather than weakening the inference. The support for P2 is conceptual: generic theism permits many intervention policies. The support for P3 is mathematical and methodological: posterior odds require prior odds and a Bayes factor, while a local explanatory ranking specifies neither. Whether a particular historical case supplies the missing inputs remains an empirical and interpretive question.

A.2 Symbols

- T : the relevant generic-theism background;
- M : the proposition that the specified miracle occurred;
- E : the public historical evidence;
- A : the judgment that M is the best explanation of E among the rivals considered;
- P : the probability function used for priors, likelihoods, and posteriors;
- q : a stated high-credence threshold, with $0.5 < q < 1$;
- C_q : the claim that $P(M \mid E, T) \geq q$;
- $B(E)$: the Bayes factor $P(E \mid M, T)/P(E \mid \neg M, T)$;
- $O_T(M)$: the prior odds $P(M \mid T)/P(\neg M \mid T)$.

A.3 Threshold consequence

Bayes' theorem yields

$$P(M \mid E, T) \geq q \iff B(E) O_T(M) \geq \frac{q}{1-q}. \quad (5)$$

Generic theism supplies neither M nor a fixed value of $O_T(M)$, so $T \not\models M$. The abductive judgment A is only a ranking among considered explanations and, without further premises, fixes neither $B(E)$ nor $O_T(M)$. Hence

$$A \wedge T \not\models C_q. \quad (6)$$

This non-entailment is deductively valid: the premises do not contain the quantities required by the threshold condition. It does not entail $\neg M$ or $P(M \mid E, T) < q$. Empirical and theological work could supply values for $B(E)$ and $O_T(M)$ that satisfy the displayed inequality. The formal conclusion is exactly the bounded central thesis: uncalibrated abduction plus generic theism does not transmit public high credence to a specific miracle.

References

- Climenhaga, N. (2017). How explanation guides confirmation. *Philosophy of Science*, 84(2), 359–368. <https://doi.org/10.1086/690723>
- Douven, I. (2017). Inference to the best explanation: What is it? and why should we care? In K. McCain & T. Poston (Eds.), *Best explanations: New essays on inference to the best explanation* (pp. 7–24). Oxford University Press. <https://doi.org/10.1093/oso/9780198746904.003.0002>
- Dretske, F. I. (1970). Epistemic operators. *The Journal of Philosophy*, 67(24), 1007–1023. <https://doi.org/10.2307/2024710>
- Earman, J. (2000). *Hume's abject failure: The argument against miracles*. Oxford University Press.
- Harman, G. H. (1965). The inference to the best explanation. *The Philosophical Review*, 74(1), 88–95. <https://doi.org/10.2307/2183532>
- Jantzen, B. C. (2012). Peirce on miracles: The failure of bayesian analysis. In J. Chandler & V. S. Harrison (Eds.), *Probability in the philosophy of religion* (pp. 27–45). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199604760.003.0002>
- Lipton, P. (2004). *Inference to the best explanation* (2nd ed.). Routledge.
- McGrew, T., & McGrew, L. (2009). The argument from miracles: A cumulative case for the resurrection of jesus of nazareth. In W. L. Craig & J. P. Moreland (Eds.), *The blackwell companion to natural theology* (pp. 593–662). Wiley-Blackwell. <https://doi.org/10.1002/9781444308334.ch11>
- Norton, J. D. (2021). *The material theory of induction*. University of Calgary Press. <https://doi.org/10.2307/j.ctv25wxcb5>
- Plantinga, A. (2000). *Warranted christian belief*. Oxford University Press. <https://doi.org/10.1093/0195131932.001.0001>
- Schupbach, J. N., & Sprenger, J. (2011). The logic of explanatory power. *Philosophy of Science*, 78(1), 105–127. <https://doi.org/10.1086/658111>
- Sober, E. (2008). *Evidence and evolution: The logic behind the science*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511806285>
- Sprenger, J., & Hartmann, S. (2019). *Bayesian philosophy of science*. Oxford University Press. <https://doi.org/10.1093/oso/9780199672110.001.0001>
- Stanford, P. K. (2006). *Exceeding our grasp: Science, history, and the problem of unconceived alternatives*. Oxford University Press.
- Swinburne, R. (2003). *The resurrection of god incarnate*. Oxford University Press. <https://doi.org/10.1093/0199257469.001.0001>