The "Unknowns Reserve" as Epistemic Norm: Bayesian Humility in Historical Miracle Claims

Phil Stilwell

Independent Scholar

The "Unknowns Reserve" as Epistemic Norm: Bayesian Humility in Historical Miracle Claims

Bayesian reasoning provides a rigorous framework for evaluating extraordinary claims, yet it is only as sound as the assumptions encoded in priors, likelihoods, and model structure. One often overlooked principle is the necessity of an "Unknowns Reserve"—a probability allocation for mechanisms not yet conceived or understood. History is replete with examples where premature attribution of phenomena to divine agency collapsed under the weight of later discoveries: lightning once attributed to Zeus yielded to atmospheric electricity, diseases once thought to be demonic afflictions yielded to microbial theory, and eclipses once feared as divine portents yielded to orbital mechanics. This paper argues that properly applied Bayesian reasoning must institutionalize such a reserve as an epistemic norm, particularly in assessing historical miracle claims that can no longer be tested under scientific scrutiny. Without this reserve, Bayesian models risk premature closure, inflating the plausibility of the miraculous and underestimating the vast space of possible natural mechanisms. By situating the Unknowns Reserve within the broader philosophy of science (Jaynes, 2003; Dawid, 2005), cognitive psychology (Loftus & Palmer, 1974; Gigerenzer, 2002), and historical epistemology (Hume, 1748/2007; Earman, 2000), the paper shows why apologetic deployments of Bayes that omit it misrepresent the mathematics. The conclusion is that the Unknowns Reserve is not skepticism for its own sake but a structural safeguard for rational inquiry in contexts where evidence is fragmentary, testimony is fallible, and explanatory space is incompletely mapped.

Opening Narrative

Bayes' theorem is often invoked as a tool for adjudicating extraordinary claims, such as the resurrection of Jesus or reports of miraculous healings (Carrier, 2012, 2014). The formalism promises neutrality: evidence is weighed against competing hypotheses, with posterior probabilities reflecting the relative strength of each. Yet the neutrality of the framework depends on whether all live alternatives are represented fairly. When unimagined or unknown mechanisms are excluded from consideration, the evidential weight is artificially funneled toward favored

hypotheses. This omission is not merely a technical oversight; it is an epistemic failure, amounting to the pretense that our models are complete when history repeatedly demonstrates otherwise (Jaynes, 2003).

The necessity of reserving probability for the unknown becomes clear in retrospect. For centuries, lightning was explained in terms of divine anger—Zeus casting bolts from Olympus. Only with the discovery of electrons and atmospheric electricity did this account give way to a naturalistic explanation. The phenomenon itself had not changed; what shifted was the range of live hypotheses considered. A Bayesian analysis conducted in the pre-scientific era without a category for "unknown natural mechanisms" would have assigned undue weight to the divine hypothesis. The lesson generalizes: where human understanding is incomplete, probability mass must be allocated to the unknown. In Bayesian terms, this means assigning a nonzero $P(U)$ to capture unmodeled alternatives, rather than implicitly fixing $P(U) = 0$.

This principle—the "Unknowns Reserve"—is not an optional courtesy but a normative requirement for rational inference. It acknowledges that our models of the world are always incomplete and that historical claims, especially those of miracles, exist in evidential contexts that are permanently underdetermined. Unlike phenomena still open to experimental or forensic scrutiny, miracle claims embedded in antiquity are closed to new direct evidence. This makes the reserve not smaller but larger. Any Bayesian analysis that omits it risks replicating the same epistemic errors that once made Zeus responsible for lightning and demons responsible for disease.

Historical Cases of Premature Closure

The case of lightning attributed to Zeus is emblematic but far from unique. Human cultures have a long record of filling explanatory gaps with divine or supernatural agency, only to see these attributions overturned once new mechanisms were discovered (Earman, 2000). Each reversal illustrates the danger of neglecting an explicit reserve for the unknown in our probability models.

Disease and Demonic Affliction

For centuries, the prevailing explanation for epidemics involved spirits or divine punishment. Medieval Europe often interpreted the plague as God's wrath, while many non-Western cultures ascribed seizures and other neurological conditions to demonic possession. In both cases, the phenomena were taken as decisive evidence for the supernatural. With the discovery of microbes and the development of germ theory, these explanations collapsed (Sober, 2008). The phenomenon of disease was always natural; what changed was the expansion of the hypothesis space.

Mental Illness and Possession

Similarly, symptoms of schizophrenia, epilepsy, and other mental illnesses were frequently understood as spiritual afflictions. The narrative of possession fit cultural expectations and carried theological weight, yet neuroscience and psychology have since traced such phenomena to cognitive and biological mechanisms (Gigerenzer, 2002). In African and Asian traditions, spirit possession continues to be invoked where modern medicine is less accessible, underscoring the persistence of this explanatory gap.

Eclipses and Divine Portents

Solar and lunar eclipses once carried immense religious and political significance. Ancient Mesopotamians viewed them as omens of dynastic change, Chinese emperors feared them as signs of cosmic disorder, and Mayan rituals framed them as moments requiring appeasement of the gods. Today, orbital mechanics provides precise and predictive models of these events, stripping them of their supposed theological import (Sober, 2008).

Comets, Earthquakes, and Other "Signs"

Comets were long feared as harbingers of disaster, earthquakes as expressions of divine anger, and famine as punishment for communal sin. In each case, the transition from supernatural attribution to natural explanation demonstrates the same pattern: the unknowns were not properly reserved, leading to epistemic overconfidence in the miraculous (Hume, 1748/2007).

General Pattern

The pattern across these cases is striking. Phenomena that were once deemed decisive evidence for divine activity are now seen as ordinary natural processes. In Bayesian terms, the failure lay not in the structure of the theorem but in the model construction: the probability mass that should have been set aside for unknowns was instead allocated to the supernatural. This lesson must inform our treatment of historical miracle claims. To exclude the Unknowns Reserve is not neutral; it is to repeat, in a new domain, the same errors that once made Zeus responsible for lightning and demons responsible for disease.

Formalizing the "Unknowns Reserve"

The Unknowns Reserve is a structural feature of Bayesian model-building that assigns explicit probability mass to mechanisms not yet specified or even imagined. All models are incomplete, and in historical contexts—especially those involving miracle claims with sparse, dependent, and non-replicable evidence—model incompleteness is not marginal but central (Jaynes, 2003; Dawid, 2005).

Definition via the Law of Total Probability

Let H denote the specific miracle hypothesis under evaluation. Let $\{H_{\text{alt},j}\}_{j=1}^{m}$ be a set of explicitly modeled non-miracle alternatives. Introduce a residual hypothesis U for unknown or unmodeled mechanisms. The law of total probability over the complement of H becomes:

$$P(\neg H) \;=\; \sum_{j=1}^{m} P(H_{\text{alt},j}) \;+\; P(U). \tag{1}$$

Annotation: The complement of H is not a single foil but a mixture of many modeled alternatives plus a residual mass for as-yet-unmodeled mechanisms U. The Unknowns Reserve is $P(U) > 0$.

Updating on evidence E yields:

$$P(H \mid E) \;=\; \frac{P(H)\,P(E \mid H)}{P(H)\,P(E \mid H) + \sum_{j=1}^{m} P(H_{\text{alt},j})\,P(E \mid H_{\text{alt},j}) + P(U)\,P(E \mid U)}. \tag{2}$$

Annotation: The denominator is the "fairness" term that allocates the evidential pie across all live sources of E. The unknowns contribution $P(U)P(E \mid U)$ prevents premature closure.

Normative Range for the Reserve

To prevent premature closure, we impose a conservative lower and upper bound for the reserve in low-information contexts:

$$0.05 \;\le\; P(U) \;\le\; 0.35. \tag{3}$$

This range reflects the fact that ancient miracle dossiers are sparse, interdependent, and structurally underdetermined.

Posterior Odds and a Normative Lower Bound for $P(U)$

The posterior odds with an explicit reserve are:

$$\frac{P(H \mid E)}{P(\neg H \mid E)} = \frac{P(H)\,P(E \mid H)}{\sum_{j=1}^{m} P(H_{\text{alt},j})\,P(E \mid H_{\text{alt},j}) + P(U)\,P(E \mid U)}. \tag{4}$$

We adopt a conservative context-dependent lower bound

$$P(U) \geq \alpha_U, \qquad \alpha_U \in [0.05,\, 0.35], \tag{5}$$

with larger α_U in older, testimony-dependent, non-instrumented dossiers.

Calibrating α_U to Context

Let R (Remoteness), D (Dependence), I (Instrumentation), C (Complexity) be scaled to $[0, 1]$. Set

$$\alpha_U = \text{clip}\big(\alpha_{\min} + \lambda_R R + \lambda_D D + \lambda_C C - \lambda_I I,\; \alpha_{\min}, \alpha_{\max}\big),$$

with $(\alpha_{\min}, \alpha_{\max}) = (0.05, 0.35)$ and λ's chosen so typical ancient miracle dossiers land near 0.15–0.30.

Monotonicity in the Reserve

Holding other terms fixed,

$$\frac{\partial P(H \mid E)}{\partial P(U)} = -\frac{P(H)\,P(E \mid H)\,P(E \mid U)}{\big(P(H)P(E \mid H) + \sum_j P(H_{\text{alt},j})P(E \mid H_{\text{alt},j}) + P(U)P(E \mid U)\big)^2} < 0.$$

Decision Thresholds

Adopt H only if $P(H \mid E) \geq \tau = \frac{c_{FP}}{c_{FP}+c_{FN}}$, with $c_{FP} \gg c_{FN}$ in public knowledge systems. Larger α_U aligns inference with conservative τ.

Worked Historical Micro-Case

Consider a stylized "resurrection dossier." Suppose three testimonies exist, but they are heavily interdependent (all derived from a single oral source). The effective sample size is therefore only $ESS = 1.5$ rather than three. Set $P(H) = 10^{-6}$, $P(E \mid H) = 0.9$, and the modeled natural alternatives (fraud, grief visions) contribute jointly 10^{-3} to the denominator in Eq. 2.

If the reserve is minimal ($\alpha_U = 0.05$ with $P(E \mid U) = 0.10$), the posterior is

$$P(H \mid E) \approx \frac{9.0 \times 10^{-7}}{9.0 \times 10^{-7} + 10^{-3} + 0.005} \approx 0.00018.$$

If instead the dossier is judged highly remote, dependent, and complex, leading to $\alpha_U = 0.25$, the denominator contribution rises to 0.025. The posterior falls further to

$$P(H \mid E) \approx \frac{9.0 \times 10^{-7}}{9.0 \times 10^{-7} + 10^{-3} + 0.025} \approx 0.000036.$$

Annotation: Even with generous priors and strong likelihood for H, the posterior remains orders of magnitude below any reasonable decision threshold. The difference between $\alpha_U = 0.05$ and $\alpha_U = 0.25$ illustrates how dossier context drives outcomes and why reserves must be explicit and reported.

What the Reserve Is Not

The Unknowns Reserve is not a "get-out-of-Bayes free card." It does not permit unlimited skepticism, nor does it guarantee that miracles are impossible. The reserve is **bounded** (see Eq. 3), **context-sensitive** (stronger in low-information dossiers, weaker where evidence is instrumented and replicable), and must be accompanied by **pre-declared sensitivity analysis**. Its purpose is not to tilt outcomes toward naturalism, but to prevent premature closure by ensuring that unmodeled mechanisms receive explicit probability mass rather than being silently excluded.

Plain-Language Summary

In simpler terms, the Unknowns Reserve is the probability we hold back for the possibility that something we have not yet thought of explains the evidence. To omit it is to assume—without

argument—that we already know all possible explanations. Apologists often make this mistake by implicitly setting $P(U) = 0$, leaving only the resurrection and a few weak natural alternatives in the model (McGrew & McGrew, 2009; Swinburne, 2003; Stilwell, 2025). This inflates the apparent probability of the miraculous. The correct approach is to leave a slice of probability for "none of the above," since history has shown again and again that previously unimaginable explanations often turn out to be true.

Unknowns in Historical Miracle Claims

The necessity of an explicit Unknowns Reserve is nowhere clearer than in the evaluation of miracle claims from antiquity. Unlike contemporary scientific hypotheses, such claims cannot be subjected to repeatable experiment, forensic re-investigation, or independent replication (Carrier, 2012). Their evidential base consists almost entirely of testimonial traditions, often written decades after the alleged events and preserved through communities deeply invested in their theological significance.

Resurrections in Antiquity

The resurrection of Jesus is the paradigmatic case, but it is not unique. Classical antiquity is full of accounts of figures such as Romulus, Asclepius, and Hercules being raised, translated, or glorified after death. Later hagiographies of Christian saints echo the same narrative patterns. Anthropological studies show that charismatic leaders are frequently elevated into divine or semi-divine figures in collective memory. Some of these mechanisms are modeled explicitly as alternatives in Bayesian analysis: legend growth, grief visions, or political invention. Yet many pathways remain unmapped. Which precise social dynamics catalyze mythic crystallization? How do oral traditions select for particular motifs? These open questions are exactly what the Unknowns Reserve represents: probability mass for natural mechanisms not yet specified.

Visions and Apparitions

Apparition traditions show the same dynamic. The Marian visions at Lourdes and Fatima, the Angel Moroni's revelations to Joseph Smith, and countless local apparitions in Hindu and Buddhist traditions all involve sincere testimony from multiple witnesses. Yet sociological

analysis demonstrates that visionary experiences often emerge in emotionally charged groups, under conditions of expectancy and contagion. While some of these mechanisms are partially understood, the pathways remain incompletely mapped. A Bayesian analysis that assigns $P(U) = 0$ ignores the vast space of psychological and cultural mechanisms that may account for such testimony.

Modern Parallels: UFOs and Paranormal Reports

Contemporary analogues further demonstrate the need for an Unknowns Reserve. UFO sightings, alien abduction reports, and paranormal claims share striking similarities with ancient miracle traditions: sincere testimony, multiple witnesses, dramatic cultural uptake, and eventual codification into narrative form. In each case, the testimonial data look, at first glance, like strong evidence for extraordinary events. Yet subsequent research uncovers mechanisms such as false memory, suggestibility, atmospheric misperception, and media amplification. Crucially, new mechanisms continue to be identified, showing that the explanatory space is still expanding. A fair Bayesian analysis must allocate probability to U to reflect these uncharted pathways.

Low-Information Evidential Contexts

A central feature of historical miracle claims is their embedding in low-information environments. Literacy was limited, records fragmentary, and alternative perspectives rarely preserved. Events were often reported in the idiom of prophecy-fulfillment or myth, where narrative and theology blended inseparably. These conditions make it impossible to enumerate all live natural explanations. The evidential record is not only incomplete but structurally underdetermined. In Bayesian terms, this sharply raises the lower bound of $P(U)$, ensuring that the denominator of (2) is not artificially narrowed.

Why the Reserve Must Grow, Not Shrink

Apologists often argue that the absence of obvious natural explanations strengthens the miraculous hypothesis. The reverse is true. In historical domains where the dossier is sparse, the testimony interdependent, and the cultural distance high, the set of unmodeled alternatives is larger, not smaller. The probability mass allocated to the unknown must therefore grow, not

contract. Just as lightning demanded a nonzero reserve before the discovery of electrons, so too must miracle claims embedded in antiquity be analyzed with substantial allocation to unmodeled causes. Any analysis that silently sets $P(U) = 0$ does not honor Bayes' theorem but subverts it.

Cognitive & Anthropological Rationale for the Reserve

Cognitive science and anthropology reinforce the need for the Unknowns Reserve. Humans are predisposed to hyperactive agency detection (HADD), making ambiguous stimuli feel purposeful (Gigerenzer, 2002). Memory is reconstructive, not archival, and can be distorted through suggestion and retelling (Loftus & Palmer, 1974). Anthropological studies document recurrent apotheosis narratives—from Romulus to Alexander the Great to the Buddha's relics—as well as the rise of Cargo Cults in Melanesia under colonial stress (Durkheim, 1912/1995). Ritual contexts generate "collective effervescence," producing intense group experiences interpreted as divine encounters (Durkheim, 1912/1995). Because these mechanisms are diverse, overlapping, and incompletely understood, they exemplify the kind of causal space that belongs to the Unknowns Reserve.

Addressing Apologetic Pushback

Objection 1: "If you cannot specify an alternative, you cannot count it."

Richard Swinburne (2003) and Timothy McGrew (McGrew & McGrew, 2009) argue that probabilities must be assigned only to explicitly described hypotheses. This objection fails to recognize that omitting a category for the unknown is itself arbitrary, amounting to the tacit assumption that $P(U) = 0$. History demonstrates the error of this assumption: natural explanations once unimaginable—microbes, tectonics, atmospheric electricity—were later discovered to account for phenomena once attributed to gods.

Objection 2: "The absence of known natural causes strengthens the miracle hypothesis."

Apologists often argue that if skeptics cannot produce a concrete natural explanation, the miracle gains in probability. Yet this inverts the logic of Bayesian modeling. Sparse evidence does not shrink the unknown space but expands it. In contexts where testimony is interdependent,

records fragmentary, and cultural distance high, the list of plausible but unarticulated mechanisms grows larger, not smaller. Thus the absence of a named natural explanation increases, rather than decreases, the probability mass that should be allocated to U. This apologetic inversion—treating the absence of a named natural mechanism as direct support for the miraculous—is the same pattern documented in detailed critiques of Bayesian resurrection apologetics (Stilwell, 2025).

Objection 3: "The reserve is a disguise for atheism or naturalism."

The Unknowns Reserve does not presuppose atheism. It does not deny the possibility of miracles; it merely acknowledges model incompleteness. Indeed, even theists appeal to analogous reserves. When confronted with contradictions or ambiguities in their theology, apologists often defer to "the mystery of God's plan" or insist that divine reasons may remain beyond human comprehension (Plantinga, 2000). Such claims are, in effect, a theistic Unknowns Reserve. The principle is therefore worldview-neutral.

Objection 4: "The reserve dilutes evidence unfairly."

Some argue that assigning probability mass to the unknown dilutes the discriminating power of evidence. This confuses caution with bias. Extraordinary claims require extraordinary evidence precisely because their priors are so small (Hume, 1748/2007). If evidence does not clearly discriminate between the miracle hypothesis and its competitors, the miracle cannot claim the evidence as uniquely its own. The Unknowns Reserve does not dilute evidence but keeps it proportionate.

Objection 5: "The reserve makes miracles unfalsifiable."

A final objection is that by always holding back probability for the unknown, skeptics insulate themselves from ever affirming a miracle. But the reserve is bounded, not infinite. In earlier sections, we proposed a normative range of 5–35% for $P(U)$ in Eq. 3, appropriate for low-information contexts. If a miracle hypothesis produced evidence so strong and so discriminative that no plausible unknown mechanism could account for it, the posterior could still climb dramatically. The reserve does not foreclose the miraculous; it simply ensures that weak evidence does not masquerade as decisive.

Symmetry of Appeal

It is worth noting the symmetry of appeals to the unknown. Apologists themselves invoke unmodeled possibilities when defending theological commitments. "God's ways are higher than ours" functions as a buffer against evidential tension. If believers may reasonably appeal to such reserves when defending doctrine, then skeptics may reasonably invoke them when defending against premature closure on miracle claims. The Unknowns Reserve is not skepticism in disguise but a principle of intellectual humility that both sides implicitly recognize.

General Lessons for Epistemology

The Unknowns Reserve extends beyond miracles. In science, it prevents premature closure in cases like dark matter, dark energy, or COVID-19 origins (Taleb, 2007). In law, it guards against wrongful convictions caused by overconfidence in eyewitness testimony (Loftus & Palmer, 1974). In medicine, it protects against misdiagnosis by leaving room for rare or novel conditions. In everyday life, it inoculates against superstition by explaining why coincidences need not be fated. Hume's dictum that extraordinary claims require extraordinary evidence (Hume, 1748/2007) can be formalized in Bayesian terms as the demand for an Unknowns Reserve large enough to account for unmodeled possibilities.

Synthesis Across Domains

Figure 1 summarizes how distinct domains— philosophy, science, psychology/anthropology, and apologetics—each converge on the necessity of an explicit Unknowns Reserve. By visualizing the diverse inputs, the diagram highlights that the reserve is not an ad hoc invention of skeptics but a principle reinforced across disciplines.

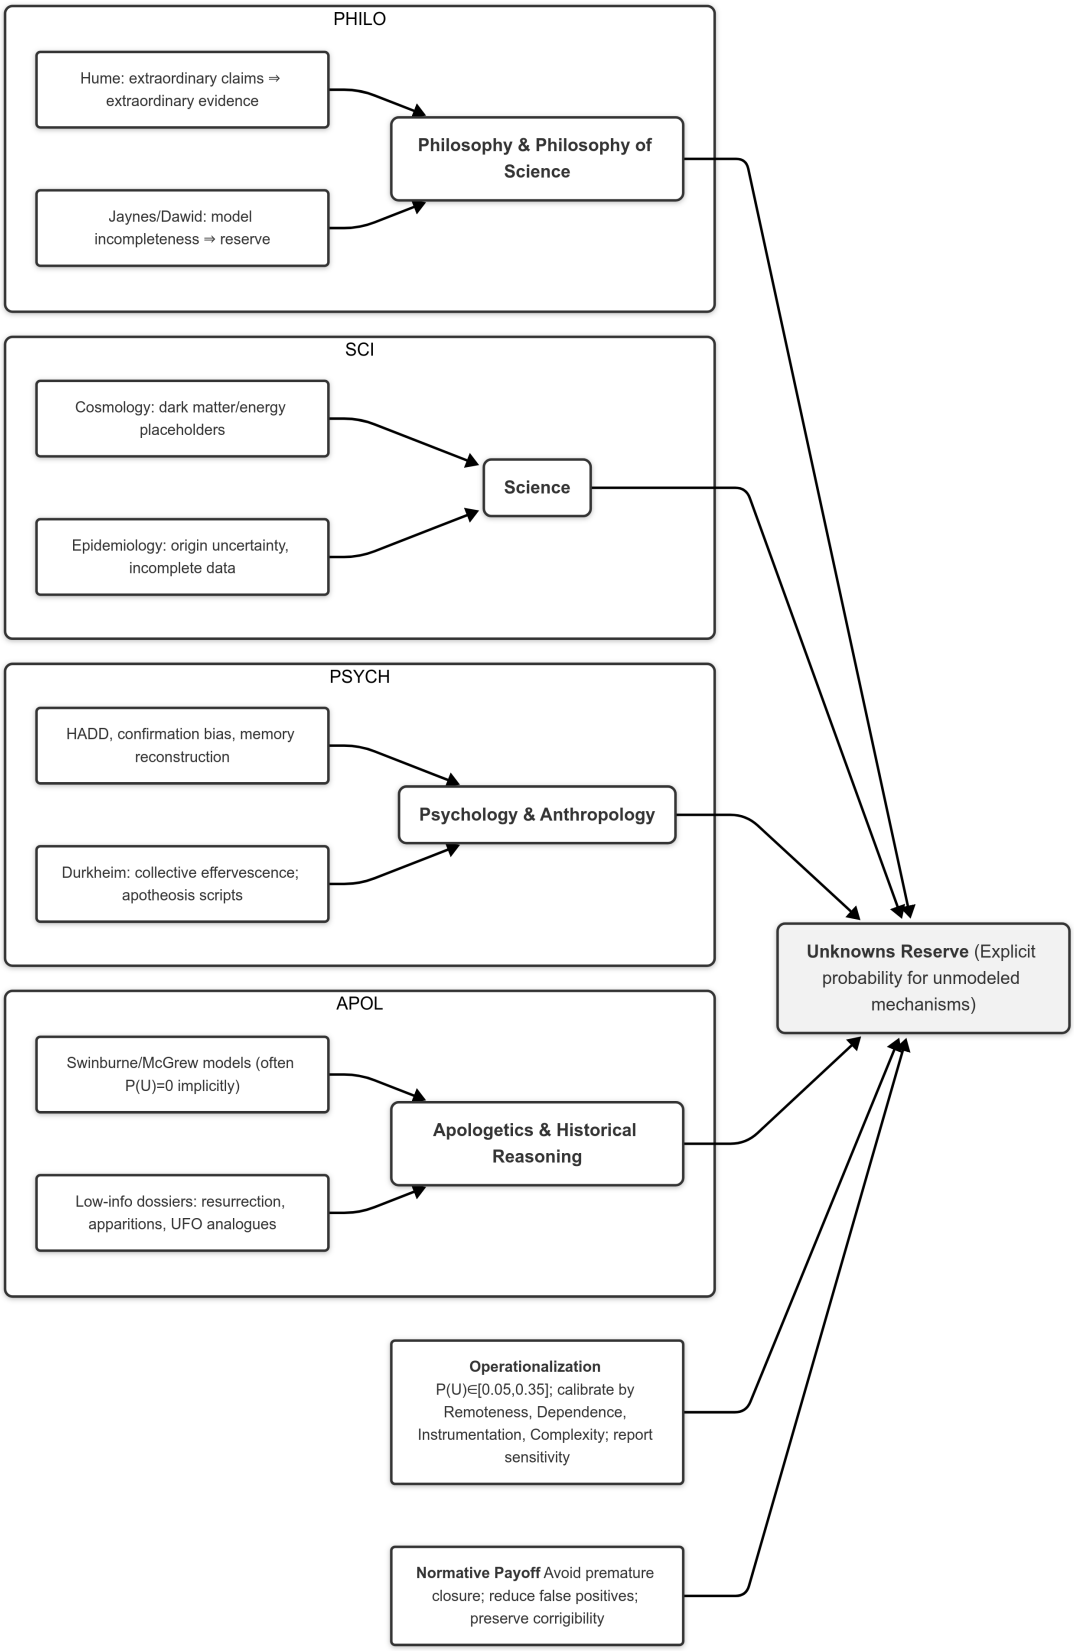

Figure 1

How multiple domains converge to justify an Unknowns Reserve in Bayesian assessments of historical miracle claims.

Conclusion

The central claim of this paper is that Bayesian reasoning requires an explicit allocation of probability mass to the unknown. Historical review shows that phenomena once seen as unambiguous signs of divine action—Zeus's lightning, demon-caused plagues, eclipses as divine portents—later yielded to natural explanations once unimaginable (Earman, 2000; Sober, 2008). Formal analysis confirms the need: omitting $P(U)$ artificially inflates the posterior for miracles. Cognitive and anthropological evidence reinforces the point: human biases, memory distortions, and ritual contagion routinely generate miracle-like testimony without extraordinary causes (Loftus & Palmer, 1974; Durkheim, 1912/1995). Addressing apologetic objections shows that the reserve is not biased but principled, and its broader application across science, law, medicine, and everyday reasoning demonstrates its universality. The irony is that Bayes' theorem, so often invoked to support miraculous certainty, when applied honestly, undermines it. The true contribution of Bayes is to remind us—with mathematical clarity—that epistemic humility is mandatory. The Unknowns Reserve is the structural embodiment of that humility.

References

Carrier, R. (2012). *Proving history: Bayes's theorem and the quest for the historical Jesus*. Prometheus Books.

Carrier, R. (2014). *On the historicity of Jesus: Why we might have reason for doubt*. Sheffield Phoenix Press.

Dawid, A. P. (2005). Probability and proof. *British Journal for the Philosophy of Science, 56*(2), 275–305. https://doi.org/10.1093/bjps/axi115

Durkheim, E. (1912/1995). *The elementary forms of religious life* (K. Fields, Trans.). Free Press.

Earman, J. (2000). *Hume's abject failure: The argument against miracles*. Oxford University Press.

Gigerenzer, G. (2002). *Adaptive thinking: Rationality in the real world*. Oxford University Press.

Hume, D. (1748/2007). *An enquiry concerning human understanding* (P. Millican, Ed.). Oxford University Press. (Original work published 1748)

Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge University Press.

Loftus, E. F., & Palmer, J. C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior, 13*(5), 585–589. https://doi.org/10.1016/S0022-5371(74)80011-3

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.

Plantinga, A. (2000). *Warranted Christian belief.* Oxford University Press.

Sober, E. (2008). *Evidence and evolution: The logic behind the science*. Cambridge University Press.

Swinburne, R. (2003). *The resurrection of God incarnate*. Oxford University Press.

Taleb, N. N. (2007). *The black swan: The impact of the highly improbable*. Random House.

Stilwell, P. (2025). *Bayesian blunders in resurrection apologetics*.