

The Inflation of Possibilities

Admissibility, Unknown Reserves, and the Discipline of Hypothesis Space

Phil Stilwell

Independent Scholar

ORCID: 0009-0009-3725-8682

Revised August 2026

Central claim

Evidence can favor one identified hypothesis over another without an exhaustive theory list, but an absolute posterior probability depends on the modeled partition. An ambiguous label supplies no single likelihood, a composite family needs an explicit treatment of its members, and probability assigned to a residual “other” category does not transfer to a favored member. This blocks unsupported evidential promotion, not metaphysical possibility or early-stage inquiry.

Preview: From possibility to warranted credit

Distinguish it	Compare it	Bound the result
Is this an ambiguous phrase, a defined family, an imprecise claim, or one stable hypothesis?	Pairwise evidence requires identified rivals; a composite family also requires a principled treatment of its members.	A restricted comparison can rank named candidates without establishing their absolute probability.
Keep unimagined alternatives in an explicit residual category.	Do not transfer the reserve’s probability to one favored member without a new evidential bridge.	Match the precision of the conclusion to the precision of the hypothesis and revise the set when checks fail.

Readiness is graduated: a conjecture may guide inquiry before it can support a qualitative comparison or a numerical posterior.

Visual preview:



The graphic separates an unknown reserve from candidates passing coherence, evidence, explanatory-cost, and—when causal effects are claimed—interface checks. Passing the filter licenses comparison; it does not guarantee truth.

Abstract

Philosophical and apologetic arguments sometimes move too quickly from a pronounceable description to a weighted alternative: “timeless cause,” “disembodied mind,” or “necessary personal explanation” is named, treated as a live rival, and then allowed to benefit from a gap in rival accounts. This paper calls that move *epistemic inflation*. The error is not entertaining an unusual possibility. It is allowing an ambiguous label or a favored member of an unknown remainder to receive support that the comparison has not supplied.

The argument distinguishes prior plausibility, incremental confirmation, and comparative favoring. It also separates four objects that are often conflated: an ambiguous label, a defined composite family, an imprecise hypothesis, and a residual complement. Pairwise Bayes factors require identified hypotheses but not an exhaustive theory list. Absolute posterior probabilities, by contrast, depend on the modeled partition; and a composite family’s likelihood depends on an explicit weighting or bounding of its members. The central non-transfer principle follows: a rise in the probability of “something else” does not by itself raise the probability of one selected member of that remainder.

The rule is graduated rather than exclusionary. A conjecture may guide research before it supports a qualitative comparison, and qualitative support may be warranted before precise numerical modeling. Historical, metaphysical, nonphysical, and theistic hypotheses need no laboratory mechanism merely to be considered. They do need stable enough content, relevant contrasts, and a defensible evidential bridge for the strength of credit claimed.

Keywords: epistemology; Bayesian confirmation; hypothesis space; unconceived alternatives; conceivability; theism; model comparison

1 The Answer Before the Framework

Human reason is generous with words. We can say “timeless cause,” “unembodied mind,” “supernatural agent,” or “infinite personal ground” before we know whether the words pick out one stable claim. A phrase can sound like an option because grammar gives it a familiar shape. The danger begins when that feeling of possibility is counted as evidence that a distinct hypothesis deserves probability or explanatory credit.

Call this mistake *epistemic inflation*: increasing the apparent supply of supported alternatives without increasing the information available for comparing them. The metaphor is intentionally polemical. It targets an inference, not the people who make it. A genuinely novel idea may begin in rough language and later become an excellent hypothesis. The complaint is that naming, comparing, and confirming are different achievements.

A rational person may have some degree of confidence in an ordinary-language claim before possessing a formal model. The present rule concerns the public support claimed from specified evidence. Before evidence can favor one candidate over another, the comparison must identify what the candidate says, how it differs from its rivals, and why the evidence is more or less expected if it is true. Before a numerical posterior is defended, the additional assumptions that turn the comparison into a probability model must also be stated.

That rule matters wherever a gap does rhetorical work. If a natural explanation is incomplete, a theistic explanation is named. If a physical account leaves consciousness difficult, a disembodied subject is added. If a historical record is fragmentary, miracle or conspiracy is placed beside ordinary causal stories. The new label may deserve investigation. It does not yet deserve the evidential weight of the entire unexplained remainder.

The conclusion is therefore answer-first: underdescription does not make a candidate false or impossible. It places a ceiling on the kind and precision of evidential credit presently warranted. A phrase may prompt inquiry; a stable qualitative hypothesis may support an ordinal comparison; a numerical posterior requires still more structure.

1.1 Possibility inflation in ordinary talk

The pattern is common because explanatory gaps naturally invite stories. Consider five familiar forms.

- **An unexplained symptom.** Standard tests have not identified a cause, so a hidden toxin, parasite, or food sensitivity is proposed. Any of these may deserve investigation. The failed tests enlarge the space of possibilities; they do not by themselves distinguish the favored diagnosis from the other untested causes.
- **A strained relationship.** A friend stops replying, and “they secretly resent me” becomes the explanation. Silence is compatible with resentment, but also with illness, distraction, embarrassment, or a lost message. The interpretation earns support only from further patterns that favor it over those rivals.
- **A political irregularity.** A reporting error or unexplained delay is described as evidence of a coordinated plot. The anomaly may warrant an audit. Promotion to conspiracy requires a stable account of who coordinated what, what traces that coordination should leave, and why those traces fit better than error or ordinary disorder.
- **A surprising computer response.** A conversational system produces an answer its user did not anticipate, and “it has become conscious” is introduced. Surprise to the user is not yet evidence that distinguishes consciousness from complex pattern generation. The claim needs content and consequences beyond impressive output.
- **A historical mystery.** The surviving record does not explain how an artifact was made, so a lost advanced civilization or nonhuman visitor is named. Ignorance about the actual method leaves room for those stories, but it does not transfer the probability of every unknown method to either one.

The proposed explanation could be true in any of these cases. Possibility inflation occurs one step later, when the absence of an accepted explanation is counted as positive evidence for one selected replacement without the comparison needed to discriminate it.

2 From Possibility to a Model

The word *possibility* performs several jobs. Much confusion disappears once they are separated.

2.1 Logical or metaphysical possibility

A claim is logically possible when it is consistent with the relevant logic. Metaphysical possibility is a stronger and disputed notion concerning what reality could have been like, not merely what a formal language permits. Establishing either can require substantial argument. David Chalmers’s careful work on conceivability distinguishes several senses of what can be imagined or known a priori before asking whether any of them support modal conclusions (Chalmers 2002). Casual verbal conceivability is much weaker than that program.

2.2 Epistemic possibility

A claim can be epistemically possible simply because a reasoner does not know it to be false. This status is cheap but useful. It marks humility. A closed box may contain a coin, a key, or something no one has guessed. Ignorance keeps options open, yet it does not distribute evidential support equally across every description someone can utter.

2.3 Ambiguous labels, composite families, and imprecise hypotheses

Three objects that may sound equally “broad” require different treatment. An *ambiguous label* shifts among incompatible meanings and therefore does not yet express one stable proposition. A *composite*

hypothesis explicitly says that one of several defined subcases is true. It is a genuine proposition, but its likelihood depends on how probability is distributed across those subcases. An *imprecise hypothesis* deliberately supplies a range of priors or likelihoods rather than one point value. It can support bounded conclusions when those ranges are principled.

This distinction matters because breadth is not itself a defect. “One of these three mechanisms produced the effect” can be a clear composite hypothesis. The defect arises when the family is left undefined, its internal weights are hidden, or the analyst selects whichever member fits the evidence after seeing it while charging the prior cost of the family only once.

2.4 A comparison-ready hypothesis

A *hypothesis* in this paper is a claim or model sufficiently definite for the inquiry to compare with relevant rivals. It need not be complete. It must, however, exclude something, preserve its identity across the comparison, and bear some defensible relation to evidence or argument. A black-box cause can qualify if it predicts a distinctive pattern even while its mechanism is unknown. A historical hypothesis can qualify through documents, timing, traces, and background rates. Experimental control is one route to evaluability, not the definition of it.

2.5 A catch-all or unknown reserve

A *catch-all* is the remainder “none of the named hypotheses is true.” An *unknown reserve* is the practical recognition that current inquiry may have missed alternatives. These devices protect against false exhaustiveness, but they do not describe one unified cause. The catch-all may include mutually incompatible explanations with radically different predictions. It is therefore a logical container before it is a likelihood-bearing theory.

Table 1 shows the resulting discipline. The categories govern what a comparison licenses; they do not rank what is ultimately real.

Table 1: Different representational objects require different evidential treatment. The table regulates evidential use, not metaphysical possibility.

Status	What is present	Model treatment	What evidence can do
Known contradiction	An inconsistent specification	Exclude from the live comparison	Nothing supports all parts together
Ambiguous label	One phrase shifting among claims	Disaggregate or defer	No single likelihood is yet defined
Composite family	A stated disjunction of subcases	Specify or bound internal weights	Evidence bears on the weighted family
Imprecise hypothesis	Stable content with probability ranges	Report bounds or sensitivity	Support conclusions no sharper than the ranges
Residual “other”	The negation of the named set	Preserve; model only with added assumptions	Prevent false closure; do not transfer its mass
Stable candidate	Content, contrasts, and consequence profile	Compare at the warranted level	Gain or lose qualitative or numerical support

3 Readiness Without Metaphysical Policing

An admissibility rule can sound like a committee deciding what reality is allowed to contain. That is not its function. Readiness is local to an inquiry and comes in degrees. An exploratory conjecture can

justify investigation. A qualitatively comparison-ready hypothesis can be judged more or less plausible than a rival. A quantitatively modeled hypothesis can support numerical likelihoods or posteriors. The governing *credit ceiling* is simple: the strength and precision of the conclusion should not exceed the strength and precision of the specification.

Three checks determine that ceiling.

1. **Content.** Is the candidate coherent enough and stable enough that the same claim remains under discussion? This is not a demand for proof of metaphysical possibility. It prevents a claim from changing after every objection.
2. **Contrast.** What does the candidate exclude, and how does it differ from the relevant rivals? A label that includes deism, classical theism, finite designers, simulation agents, process views, and undefined transcendence is not one hypothesis unless it is explicitly treated as a family.
3. **Consequence.** What evidence or argument bears differently on the candidate and its rivals? A hypothesis can meet this check through prediction, retrodiction, explanatory constraint, conceptual entailment, testimony, or a distinctive pattern. The route need not be experimental.

A fourth check, *interface*, is conditional. When a proposed cause explains physical, temporal, or historical effects, the account must say enough about its connection to those effects to generate consequences. This need not be a microscopic mechanism. The orbit of an unseen planet was evidentially accessible before anyone observed the planet directly because the hypothesis constrained visible motion. A causal posit that can accompany every outcome without changing expectation supplies no comparable leverage.

The language of support must also be kept straight. A *prior* expresses plausibility before the stated evidence is added. *Incremental confirmation* occurs when the evidence raises a hypothesis's probability. *Comparative favoring* occurs when the evidence shifts the odds toward one candidate rather than another. An independent argument may alter the prior or provide additional evidence, but it should not be folded into the claim that this particular evidence favors the hypothesis.

Explanatory cost matters once a candidate is clear enough to assess. Simplicity, auxiliary assumptions, adjustable parameters, and fit with background knowledge affect how much a candidate explains rather than merely accommodates. Elliott Sober's analyses of parsimony emphasize that simpler is not automatically truer; the relevance of simplicity depends on how it connects to evidence and the competing models (Sober 2008; Sober 2015). Cost cannot replace content, but neither should a flexible candidate receive unlimited adjustment for free.

4 Pairwise Evidence and Absolute Probability

Bayesian reasoning sharpens the distinction between a restricted comparison and an absolute posterior. For two identified hypotheses, the Bayes factor supplied by evidence E is

$$BF_{12} = \frac{P(E | H_1)}{P(E | H_2)}. \quad (1)$$

It updates prior odds into posterior odds:

$$\frac{P(H_1 | E)}{P(H_2 | E)} = BF_{12} \frac{P(H_1)}{P(H_2)}. \quad (2)$$

This pairwise result does not require H_1 and H_2 to exhaust every possibility. It says how E shifts their relative standing. Robert Kass and Adrian Raftery emphasize this odds-based role of Bayes factors

(Kass and Raftery 1995). The comparison still requires stable hypotheses, a clear evidence statement, and defensible likelihoods, but no catch-all appears in the ratio.

Absolute posterior probability is different. For mutually exclusive and exhaustive modeled hypotheses $H_1, \dots, H_n$,

$$P(H_i | E) = \frac{P(E | H_i)P(H_i)}{\sum_{j=1}^n P(E | H_j)P(H_j)}. \quad (3)$$

The denominator normalizes across the full modeled partition. The formula does not choose that partition, individuate its members, or supply their priors and likelihoods. A result computed only over H_1 and H_2 can therefore be reported honestly as $P(H_1 | E, H_1 \vee H_2)$. It should not silently become $P(H_1 | E)$.

Composite hypotheses add a further layer. If C is partitioned into mutually exclusive submodels $M_1, \dots, M_m$, with internal weights $w_k = P(M_k | C)$, then

$$P(E | C) = \sum_{k=1}^m P(E | M_k)w_k. \quad (4)$$

The family is not meaningless; its likelihood is a weighted average. But the weights do real work. A broad family cannot use whichever member fits after the evidence arrives while paying the prior cost of only one sharply predictive member. Katie Steele and Charlotte Werndl show why simple bans on using calibrating data for confirmation are too crude: Bayesian and model-selection judgments depend on how a family, its parameters, and its calibration procedure are represented (Steele and Werndl 2018). The present rule is accordingly narrower. It forbids hidden selection and hidden weighting, not every confirmation claim involving an adjusted model.

John Earman's review of Bayesian confirmation stresses both the power of the framework and the unresolved choices involved in applying it (Earman 1992). Andrew Gelman and Cosma Shalizi likewise argue that model checking and revision do work that posterior updating alone cannot perform (Gelman and Shalizi 2013). Probability can discipline a comparison only when the comparison remains answerable to the possibility that its models are wrong.

5 The Non-Transfer Principle

The paper's central inference is not the trivial claim that only comparable things can be compared. It is the *non-transfer principle*: probability assigned to a disjunction does not automatically belong to a favored member of that disjunction.

Suppose, only for illustration, that an analyst models three exhaustive options H_1 , H_2 , and a residual R . The priors are .45, .45, and .10, while the evidence likelihoods are .10, .10, and .90. The unnormalized posterior weights are therefore .045, .045, and .090, giving posteriors .25, .25, and .50. The example assumes that the analyst has supplied a meaningful value for $P(E | R)$; that extra modeling assumption is already doing substantial work.

Even under that assumption, the conclusion is only that the residual category now has probability .50. If R is partitioned into incompatible possibilities $K_1, K_2, \dots$, then

$$P(R | E) = \sum_k P(K_k | E). \quad (5)$$

The sum does not determine any summand. A newly proposed K_1 cannot inherit .50 merely because it belongs inside R or happens to fit E . Its share depends on its prior position within the reserve and on

how E bears on it relative to the other members. The same point holds without numbers: support for “one of the named accounts is wrong” is not support for whichever replacement a speaker prefers.

6 Unknown Reserves and Unconceived Alternatives

Scientific history supplies repeated reminders that the best current set may omit the truth. P. Kyle Stanford develops this as the problem of unconceived alternatives: earlier investigators often failed to imagine theories that later matched the evidence better than the available choices (Stanford 2006). His later analysis of the catch-all problem argues that pairwise comparisons can remain possible even when the absolute probability of the winner cannot be established (Stanford 2009). The lesson is not that every unimagined claim deserves equal prior weight. It is that confidence in a winner should be sensitive to confidence that the candidate set is adequate.

One response adds a catch-all $R = \neg(H_1 \vee \dots \vee H_n)$. This makes the partition logically exhaustive, but the complement alone supplies no probability distribution over its members. A likelihood for R can be introduced through a mixture model, probability bounds, or other assumptions. Those additions may be useful, but they are not contained in the bare phrase “none of the above.” Their influence should therefore be visible in the conclusion and tested for sensitivity.

Olav Benjamin Vassend describes this as an M-open case and reviews two common strategies. One adds a catch-all despite the difficulty of assigning it a nonarbitrary prior and likelihood. Another treats the probabilities among named candidates as *suppositional*: they are conditional on one of the considered hypotheses being correct (Vassend 2026). The second strategy is especially clear when reported honestly. “Among these three models, given that one is right, H_2 performs best” is weaker than “ H_2 is probably true.”

The unconceived-alternatives challenge is not uncontested. Matthias Egg argues that well-supported causal knowledge can justify selective realism even when complete alternative spaces are unavailable (Egg 2016). This reply identifies an important asymmetry: stable interventions, multiple routes of detection, or other causal constraints can make some claims robust across rival theories. The present argument should welcome that result. A causal bridge is precisely the kind of extra structure that can raise a claim above a verbal possibility. What it cannot do is make an unspecified remainder automatically favor one particular explanation.

The unknown reserve should therefore perform three jobs:

- prevent the named list from masquerading as exhaustive;
- keep conclusions conditional on the adequacy of that list; and
- trigger model criticism when all named candidates predict poorly.

It should not perform a fourth job: secretly becoming the posterior of a preferred theory. “Something else” is not equivalent to “this particular God-concept,” “an immaterial subject,” or “a miracle.” Moving reserve weight into one of those claims requires an additional argument that individuates the claim and connects it to the evidence.

7 Conceivability as a Prompt, Not a Prior

Verbal conceivability is psychologically inexpensive. A person can repeat “timeless choice” or “nonspatial perception” without immediately encountering a contradiction. That achievement may justify inquiry. It does not establish that one stable model has been supplied.

This point should not be confused with the crude maxim that whatever cannot be pictured is impossible. Mathematics, physics, and philosophy routinely outstrip everyday imagination. Nor does the point dismiss sophisticated conceivability arguments. Chalmers distinguishes positive from negative conceivability,

primary from secondary possibility, and ideal from merely apparent conceivability precisely because the bridge from thought to modality is demanding (Chalmers 2002). The casual move at issue here does none of that work.

Three questions reveal whether conceivability has matured into a candidate:

1. What would make the description false rather than merely redescribed?
2. Which nearby claims does it exclude?
3. What feature of the evidence or argument becomes less surprising if it is true?

A reasoner need not answer each question numerically. The answers must be clear enough to prevent a phrase from changing content as the comparison proceeds. Until then, conceivability is an invitation to specify, not a voucher for evidential credit.

8 A Fair Test: Bayesian Theism

Theism is the general position that at least one deity exists. A *God-concept* is a specified conception of such a deity. This distinction matters because the bare label can cover claims with different attributes, actions, and evidential consequences.

Serious natural theology does not always stop at a label. Richard Swinburne, for example, defines a detailed theistic hypothesis, argues that its intrinsic simplicity supports its prior probability, and assesses how evidence concerning a law-governed universe, embodied persons, religious experience, and evil bears on it (Swinburne 2004). Whether that case succeeds is not settled by the present admission rule. Swinburne's proposal has moved beyond "perhaps something supernatural" and must be engaged on its actual priors, likelihoods, auxiliary assumptions, and rivals.

Consider a compact comparison. Let T be one specified theistic model, N one specified naturalistic model, R the residual category, and E a law-governed universe containing embodied conscious agents. A defender of T might argue that E is more expected under T than under N , so that $BF_{TN} > 1$. If the likelihoods are justified, that result genuinely favors T over N . The admission rule does not block it.

The result nevertheless has limits. The prior odds between T and N still matter. The likelihood under T depends on claims about the deity's reasons and possible actions; the likelihood under N depends on the natural processes and selection effects included in that model. Most importantly, $BF_{TN} > 1$ does not by itself establish $P(T | E)$, because alternatives inside R may also expect E . A restricted comparison remains useful when its restriction remains visible.

This example clarifies what the rule does and does not accomplish. First, specification can earn readiness without earning a high probability. Second, simplicity is contestable. Calling a God-concept simple does not determine whether unlimited powers, intentions, moral attributes, and freedom compress the evidence or create a large family of possible actions. Third, compatibility is not prediction. If the deity could have sufficient reason for almost any possible world, observed order may fit the hypothesis without being sharply expected by it. Fourth, rivals must be generous. A theistic hypothesis should not be compared only with one crude naturalistic story, and critics should not compare it only with an undefined negation.

The same fairness applies outside theology. A physicalist cannot dismiss an interactionist mind merely for lacking today's mechanism if it predicts a distinctive, stable pattern. A skeptic cannot dismiss a historical anomaly merely because it is rare. Admission discipline is not ideological border control. It is the demand that every side pay for evidential credit with content and comparative work.

9 Strong Objections

9.1 “Credence does not wait for a scientific model”

Correct. People rationally hold graded beliefs about elections, personal motives, historical events, and metaphysical claims without formal likelihood functions. The paper does not prohibit such credences. It denies a narrower inference: an ambiguous phrase cannot acquire a well-defined *separate* likelihood merely because a speaker assigns it confidence. Personal confidence may be real while the public evidential bridge remains underdetermined.

This distinction also prevents dogmatism. Assigning literal probability zero to a genuine possibility can make ordinary conditionalization unable to recover. Deferral is not zero. It says that the current model does not yet support a separate numerical entry at the claimed level of precision.

9.2 “Mechanisms are often unknown”

Also correct. Neptune, genes, pathogens, and historical authors were often inferred before their mechanisms or identities were fully known. What made the hypotheses serious was not complete causal machinery but constraint: they organized evidence, excluded alternatives, and risked failure. The interface check asks for a defensible consequence profile when causal effects are claimed, not for a miniature physical diagram.

Some explanations are intentionally high level. Intentions can explain action, market structure can explain prices, and selection can explain trait distributions without being replaced by microphysics. The admission test is pluralistic about explanatory levels. It asks what difference the posit makes at its own level.

9.3 “The standard is hostile to early theories”

Only if the credit ceiling is mistaken for a single gate. A sketch can enter inquiry when it makes one stable commitment. “There is an unseen body whose gravity perturbs this orbit” is incomplete but risky. “An unseen cause produced whatever orbit occurred” is not. Exploratory conjecture, qualitative comparison, and quantitative modeling are different levels of readiness. A research lead needs less specification than a numerical posterior or an assertion of best explanation.

9.4 “The catch-all is just as vague”

The bare complement is. That is why the unknown reserve should not be treated as an ordinary competitor with a confidently known likelihood. Its role can remain metacognitive: it records possible misspecification and limits the force of comparisons among named models. Alternatively, an analyst may impose a mixture distribution, use probability bounds, or test several residual models. Those methods can make the reserve mathematically useful, but the resulting conclusions depend on the added assumptions. Ignorance alone does not supply one predictive distribution.

9.5 “A model fitted to the evidence can still be confirmed”

Agreed. The crude rule “data used to build a model can never confirm it” fails across important Bayesian and classical settings. Calibration may change the relative support of a base model, and a composite model can be evaluated by averaging over its submodels. The objection therefore defeats a blanket novelty requirement, not the narrower discipline defended here. The analyst must represent the selection or calibration process, use the whole family rather than only its best-fitting member, and avoid presenting conditional fit as unconditional confirmation. Steele and Werndl’s analysis makes this qualification especially important (Steele and Werndl 2018).

9.6 “Explanatory virtues can support a theory before precise predictions”

They can. Simplicity, unification, fit with background knowledge, and causal intelligibility may provide reasons without a novel numerical prediction. Their force still depends on discriminating content. A

theory unifies only what it actually derives; it is simple only under a relevant representation; and causal intelligibility matters only when it connects the proposed cause to the effect better than its rivals do. These virtues can justify qualitative or comparative support even when point likelihoods are unavailable. What they cannot justify is an evidential claim more precise than the comparison that gives the virtues meaning.

10 Safeguards for Hypothesis Space

The practical discipline can be stated as eight questions.

1. **What exactly is the claim?** Replace an umbrella label with the proposition or model actually assessed.
2. **What kind of object is it?** Distinguish an ambiguous label, composite family, imprecise hypothesis, stable candidate, and residual complement.
3. **What does it exclude?** Identify nearby variants that would produce a different comparison.
4. **What is the rival set?** State whether the analysis is exhaustive, restricted, or conditional on one named candidate being correct.
5. **What kind of support is claimed?** Separate prior plausibility, incremental confirmation, comparative favoring, and absolute posterior probability.
6. **Why these weights and likelihoods?** Give reasons, ranges, or sensitivity analyses rather than converting familiarity or post-hoc fit into weight.
7. **Where are the unknowns?** Preserve a residual possibility of misspecification without assigning its probability to a favorite.
8. **What would force revision?** Name evidence, argument, or model failure that would lower confidence or change the candidate set.

These questions expose two common asymmetries. The first makes a favored hypothesis precise enough to collect confirming evidence but flexible enough to explain away adverse evidence. The second demands a complete mechanism from rivals while allowing the favored cause to remain a label. A fair comparison uses the same admission standard in every direction.

The order of inquiry is therefore not a rigid pipeline but a recurring loop: formulate, compare, check, revise. Priors and likelihoods can be provisional. New evidence can reveal that the candidates were badly drawn. Model checking can return an admitted hypothesis to the workshop. What discipline forbids is skipping directly from a word to a victory.

11 Conclusion

Epistemic inflation grants support because a phrase is pronounceable, unrefuted, or names an unexplained remainder. Four distinctions block it: ambiguous labels lack stable content; composite families need explicit internal treatment; pairwise evidence need not exhaust every possibility; and absolute posteriors depend on modeled partitions. Above all, probability assigned to a residual “other” does not transfer to a favored member without a further evidential bridge.

This conclusion is bounded. Historical, metaphysical, nonphysical, and theistic hypotheses can earn support without laboratory mechanisms. Conjectures may guide inquiry, specified hypotheses may earn qualitative support, and fully modeled hypotheses may justify numerical conclusions. Unknowns deserve room; named claims deserve no more precision or probability than their content and evidential bridges earn.

A Formal Reconstruction of the Central Thesis

The thesis contains two arguments: one about the content required for separate evidential credit, and one about the transfer of probability from an unknown remainder.

Argument P: Individuation and Evidential Credit

- P1.** A claim that evidence E supports a candidate at a stated level is warranted only if stable content and a defensible relation between E and relevant alternatives have been supplied.
- P2.** A merely pronounceable or ambiguous label identifies no single candidate, contrast class, or consequence profile.
- P3.** Being unrefuted, casually conceivable, or compatible with an unexplained remainder supplies no additional bridge from E to that label.
- P4. Conclusion:** On that information alone, the label is not warranted separate evidential credit. It should be clarified, disaggregated, or kept as an exploratory conjecture.

Prior plausibility or an independent argument may support a later, better-defined claim, but neither makes the original evidence relation determinate by stipulation.

Argument Q: Non-Transfer from a Residual Category

- Q1.** If residual category R is partitioned into mutually exclusive alternatives $K_1, K_2, \dots$, its posterior is the sum of their posteriors.
- Q2.** Knowing the value of a sum does not determine the value of any one term in that sum.
- Q3.** Naming favored member K_f after learning the evidence supplies neither its prior share of R nor its likelihood relative to the other members.
- Q4. Conclusion:** An increase in $P(R | E)$ cannot therefore be transferred to $P(K_f | E)$ without an additional model or argument.

Q4 allows K_f to earn later support; it blocks the move from “the named models may be wrong” to “this replacement is supported.”

Validity, Support, and Scope

Both inferences are valid at their stated strength. Argument P is a warranted-claim argument: P1 states the minimum evidential bridge, and P2–P3 establish that it has not been supplied. P1 is methodological rather than a probability theorem; it keeps the proposition, comparator, and evidence relation stable.

Argument Q rests on elementary probability once R is partitioned into mutually exclusive members. Their total does not fix any member’s share. A defender can defeat the application by supplying the missing internal weights and likelihoods; doing so confirms rather than evades the principle.

Neither argument establishes metaphysical impossibility, requires experimentation, assigns probability zero, or rules out qualitative support. Together they establish a credit ceiling: evidential conclusions may not outrun the specification and support relation supplied.

References

- Chalmers, David J. (2002). “Does Conceivability Entail Possibility?” In: *Conceivability and Possibility*. Ed. by Tamar Szabó Gendler and John Hawthorne. Oxford: Oxford University Press, pp. 145–200. DOI: [10.1093/oso/9780198250890.003.0004](https://doi.org/10.1093/oso/9780198250890.003.0004).
- Earman, John (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.
- Egg, Matthias (2016). “Expanding Our Grasp: Causal Knowledge and the Problem of Unconceived Alternatives.” In: *The British Journal for the Philosophy of Science* 67.1, pp. 115–141. DOI: [10.1093/bjps/axu025](https://doi.org/10.1093/bjps/axu025).
- Gelman, Andrew and Cosma Rohilla Shalizi (2013). “Philosophy and the Practice of Bayesian Statistics.” In: *British Journal of Mathematical and Statistical Psychology* 66.1, pp. 8–38. DOI: [10.1111/j.2044-8317.2011.02037.x](https://doi.org/10.1111/j.2044-8317.2011.02037.x).
- Kass, Robert E. and Adrian E. Raftery (1995). “Bayes Factors.” In: *Journal of the American Statistical Association* 90.430, pp. 773–795. DOI: [10.1080/01621459.1995.10476572](https://doi.org/10.1080/01621459.1995.10476572).
- Sober, Elliott (2008). *Evidence and Evolution: The Logic Behind the Science*. Cambridge: Cambridge University Press. DOI: [10.1017/CB09780511806285](https://doi.org/10.1017/CB09780511806285).
- (2015). *Ockham’s Razors: A User’s Manual*. Cambridge: Cambridge University Press. DOI: [10.1017/CB09781107705937](https://doi.org/10.1017/CB09781107705937).
- Stanford, P. Kyle (2006). *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press. DOI: [10.1093/0195174089.001.0001](https://doi.org/10.1093/0195174089.001.0001).
- (2009). “Scientific Realism, the Atomic Theory, and the Catch-All Hypothesis: Can We Test Fundamental Theories Against All Serious Alternatives?” In: *The British Journal for the Philosophy of Science* 60.2, pp. 253–269. DOI: [10.1093/bjps/axp003](https://doi.org/10.1093/bjps/axp003).
- Steele, Katie and Charlotte Werndl (2018). “Model-Selection Theory: The Need for a More Nuanced Picture of Use-Novelty and Double-Counting.” In: *The British Journal for the Philosophy of Science* 69.2, pp. 351–375. DOI: [10.1093/bjps/axw024](https://doi.org/10.1093/bjps/axw024).
- Swinburne, Richard (2004). *The Existence of God*. 2nd ed. Oxford: Oxford University Press. DOI: [10.1093/acprof:oso/9780199271672.001.0001](https://doi.org/10.1093/acprof:oso/9780199271672.001.0001).
- Vassend, Olav Benjamin (2026). “On the Ecological and Internal Rationality of Bayesian Conditionalization and Other Belief Updating Strategies.” In: *The British Journal for the Philosophy of Science* 77.1, pp. 219–243. DOI: [10.1086/724447](https://doi.org/10.1086/724447).