

The Grammar of Knowing

What Conscious Knowers Actually Have

CF Dietz

Abstract

This paper takes no position on whether truth exists as a property of propositions. We make an epistemological claim, not an ontological one: Justified True Belief is the wrong framework for describing what conscious knowers actually have when they know, because it cannot distinguish between truth and maximum probable belief within the best available grammar. That distinction is not semantic. It is the distinction on which the entire history of scientific progress turns, and JTB has no vocabulary for it. We identify three structural features of conscious knowing that jointly characterize the epistemic condition JTB cannot describe: the grammar-entry of belief formation, the directional drift of linguistic transmission, and the grammar-relativity of all knowledge claims. We define grammar precisely as a stabilized closure regime whose identity criteria constitute facts within its scope, generate remainder at its boundary, and are subject to supersession when remainder accumulates sufficient force. This framework extends Kuhn's account of paradigm change by formalizing the mechanism he identified intuitively and supplying the epistemological consequence his account left underspecified. The Copernican revolution demonstrates the structure: the Ptolemaic astronomers satisfied every condition JTB specifies, yet JTB can only say their belief was false, a verdict that does no explanatory work and misdescribes what actually happened. We then demonstrate the same structure in compound clinical knowledge claims, where a specific condition we call inferential absence renders compound claims formally unevaluable: not uncertain, not low-probability, but lacking any basis for probability assignment at all. We show that this condition exceeds the reach of imprecise probability frameworks, which require some evidential basis for bounding intervals, and of transportability and partial identification frameworks, which require that some structure connecting trial conditions to target conditions has been established. In place of JTB we propose Justified Probable Belief as a framework that accurately describes the epistemic condition conscious knowers actually occupy. We present JPB as a framework proposal, not a completed formal theory, and identify explicitly what a completion would require.

1. The Question JTB Cannot Answer

The standard account of knowledge is deceptively simple. A person S knows that P if and only if P is true, S believes that P, and S is justified in believing that P. The account descends from Plato's effort in the *Meno* and *Theaetetus* to distinguish knowledge from mere belief by requiring that belief be anchored to something stable. Its modern formulation is crisp enough to have generated a literature of enormous productivity and a problem of corresponding depth.

Edmund Gettier showed in 1963 that the three conditions are not jointly sufficient. Cases can be constructed in which all three are satisfied and the subject nonetheless does not intuitively seem to know. The philosophical literature since Gettier has largely devoted itself to identifying

what additional condition would rescue the account: causal conditions, sensitivity conditions, safety conditions, no-false-lemmas conditions, proper-function conditions. Each repair addresses Gettier's specific problem. None addresses the problem this paper identifies.

We begin with a clarification of target. This paper does not argue that JTB fails because conscious knowers cannot recognize when the truth condition is satisfied. Contemporary epistemology has long acknowledged that knowledge need not be self-certifying or transparent to the knower. Goldman's reliabilist account grounds knowledge in reliable belief-forming processes without requiring that the knower have reflective access to those processes (Goldman, 1979). Williamson's knowledge-first framework makes knowledge the fundamental epistemic state from which justification and other notions are derived, explicitly denying that knowledge must be luminous or self-presenting (Williamson, 2000). Anti-luminosity arguments establish that even one's own mental states need not be always knowable from the inside. We are not attacking the transparency assumption because JTB does not require transparency, and neither do its most sophisticated successors.

However, neither Goldman's reliabilism nor Williamson's knowledge-first account can make the distinction our argument requires. Goldman grounds knowledge in reliable belief-forming processes. But a reliable process operating within a superseded grammar is not distinguishable, by Goldman's criterion, from a reliable process operating within the current best grammar. The Ptolemaic astronomer's belief-forming processes were reliable within their grammar: they consistently produced beliefs that matched observations, generated successful predictions, and survived critical scrutiny. Reliabilism would certify their beliefs as knowledge. What reliabilism cannot explain is what changed epistemically when the grammar revised, because the revision was not driven by a failure of reliability. It was driven by the accumulation of remainder. Williamson's knowledge-first framework faces the same limitation: it makes knowledge primitive but does not provide resources to distinguish knowledge within a superseded grammar from knowledge within the current best grammar, because neither framework has a concept of grammar in the sense we have defined it. Our target is not the transparency assumption these accounts abandon. Our target is the grammar-independence assumption they retain.

A deeper objection must be addressed directly. The referee-level response to this argument is: JTB already marks the distinction you say it cannot mark. Justification tracks what is maximally warranted from within a framework. Truth is the external condition that makes knowledge knowledge. So the grammar-relativity of justification is already accommodated within JTB, and your argument establishes only a problem for first-person epistemic description, not for the factive analysis of knowledge itself.

This objection assumes that marking a distinction and doing explanatory work with it are the same thing. They are not. Truth in JTB is supposed to do three specific pieces of explanatory work. First, it anchors the difference between knowledge and lucky belief: the believer who happens to be right has truth, the believer who has justified false belief does not. Second, it explains the value of knowledge over mere justified belief: knowing that P is better than merely being justified in believing P because knowledge is factive and justified belief is not. Third, it provides the criterion for error correction: we identify error by discovering that a belief is false and replace it with a true one. The grammar-relativity argument defeats all three pieces of explanatory work, not merely the first-person access to them.

On the first: if truth is grammar-independent but the knower's best epistemic position is always grammar-relative, truth cannot anchor the distinction between knowledge and lucky belief in any way a knower or community can use, apply, or reason with. The anchor is permanently submerged. On the second: if truth is indistinguishable from maximum probable belief within the best available grammar at any point in inquiry, it cannot explain why knowledge is more valuable than justified belief, because no one can tell the difference between having it and not having it. The value claim becomes vacuous. On the third: the Copernican case shows that what actually happened was not error correction in JTB's sense. The geocentric belief was not discovered to be false and replaced with a true belief about the same proposition. A new grammar was constructed in which the proposition itself changed its character. JTB's response, that the geocentric belief was simply false, is available but does no explanatory work. It does not explain why the belief was the right thing to have, how progress happened without error in the ordinary sense, or what the epistemic achievement of the revision consisted in. JPB explains all three. That is the difference between marking a distinction and doing explanatory work with it.

Our target is different and more fundamental. We argue that JTB is the wrong framework for describing what conscious knowers actually have when they know, because it cannot distinguish between two conditions that any adequate epistemology must be able to distinguish: truth and maximum probable belief within the best available grammar. That distinction is not semantic. It is the distinction on which the entire history of scientific progress turns. JTB has no vocabulary for it. A framework that cannot make that distinction cannot describe what science does, what scientific progress means, or what was lost and gained in any episode of theory change. That is not a peripheral deficiency. It is a failure at the center of what epistemology is for.

A clarification of scope is required here. We are not arguing that JTB must serve as a complete theory of scientific progress, communication, or evidence translation. Those are separate projects with their own literatures. Our argument is narrower and more precise: JTB is a proposed analysis of what conscious knowers have when they know, and it fails at that specific task because the epistemic condition of conscious knowers is grammar-relative in a way JTB cannot accommodate. The failure is not that JTB omits a theory of science communication. The failure is that JTB cannot describe the epistemic condition of the Ptolemaic astronomer, the modern physicist, or the clinician reading a trial result, because all of them are grammar-inhabiting knowers whose best epistemic achievement is maximum probable belief within their best available grammar, and JTB has no way to describe that condition or distinguish it from what it calls truth. That is a failure of JTB as an analysis of knowledge, not a complaint that JTB fails to be a theory of everything.

We take no position in this paper on whether truth exists as a property of propositions. That is a metaphysical question we bracket explicitly, in the tradition of epistemological projects that proceed without settling the underlying ontology. Our claim is epistemological: JTB's truth condition, whatever its metaphysical status, cannot serve as a description of the epistemic achievement conscious knowers make, because the condition they actually occupy is grammar-relative in a way that JTB cannot accommodate and cannot distinguish from what it calls truth.

One further clarification is required before proceeding, and it concerns a genuine tension in the framework that a careful reader will notice. We claim to bracket metaphysics, yet the engine of our argument, remainder, depends on a specific relationship between finite grammars and the world. Remainder is the structural mismatch between what a grammar can model and what the world contains. That claim requires a minimal ontological commitment: that the world is not

exhausted by any finite grammar brought to bear on it. Something always exceeds what any closure regime can model. We cannot entirely avoid that claim without abandoning the concept of remainder entirely.

We accept this commitment and state it explicitly. It is weaker than full realism and compatible with most anti-realist positions. We do not claim to know what the world contains beyond any grammar. We do not claim there is a grammar-independent description of what the world ultimately is. We claim only that whatever the world is, it is not exhausted by any finite set of identity criteria, distinctions, and lawful relationships. That minimal commitment is what the concept of remainder requires, and it is all it requires. A framework that makes only this minimal ontological claim while remaining neutral on the larger questions of realism and anti-realism is not metaphysically evasive. It is metaphysically precise. The floor we stand on is as low as it can be while still supporting the argument. Future metaphysical grammars may have something to say about even this minimal commitment. We acknowledge that possibility: it is itself an instance of remainder.

One further clarification is required before proceeding. When we say that facts as storable, reasoned-about, and actionable are constituted within grammars, we are making an epistemological claim about the conditions under which facts become available for belief, assertion, and inference. We are not making the stronger ontological claim that there is no mind-independent world or that reality is constructed by conceptual schemes. Whether there is a world beyond any grammar brought to bear on it is a question we do not answer. What we argue is that facts as we can state them, reason about them, and act on them are always constituted within some grammar, and that this constitutive relationship has consequences for how we understand knowledge that JTB cannot accommodate. This position is epistemological constructivism about knowable facts, not metaphysical constructivism about reality. The distinction is not merely diplomatic. It is the difference between a claim about the conditions of cognition and a claim about the structure of being.

We define grammar precisely here so that the term carries consistent meaning throughout. A grammar is a stabilized closure regime: a structured set of identity criteria, distinctions, and lawful relationships that constitute facts within its scope, generate remainder at its boundary, and are subject to supersession when remainder accumulates sufficient force. Identity criteria determine what counts as the same entity across observations. Distinctions determine what differences are legible within the grammar. Lawful relationships determine what follows from what. Without a grammar in this sense there are no storable facts, only undifferentiated experience. Facts are not discovered within grammars. They are constituted by them. Every grammar is finite: it makes some distinctions and leaves others undrawn, makes some relationships legible and leaves others invisible. The structural mismatch between what a grammar can model and what the world contains is its remainder. Remainder accumulates within any successful grammar until a new grammar becomes possible, one that reorganizes the same observations within different identity criteria, dissolving the remainder the old grammar had accumulated.

This definition grounds the falsifiability of the framework. A grammar change is identifiable by the dissolution of accumulated remainder, not retroactively stipulated for any convenient theory change. Grammar supersession requires that the new grammar explain what the old grammar absorbed but could not resolve. That criterion is empirically assessable and distinguishes genuine grammar change from mere theoretical revision within a single grammar.

A clarification about the strength of this claim is essential to avoid a persistent misreading. The claim that facts as storable, reasoned-about, and actionable are constituted within grammars is the weak version of a family of claims about the relationship between mind and world. We do not make the stronger claim that facts as such are constructed by minds, or that there is no mind-independent world, or that reality is exhausted by what any grammar can model. The minimal ontological commitment stated in the preceding paragraphs is all the claim requires. What we argue is more precise: any fact that could figure in a knowledge claim, any fact to which JTB's truth condition could coherently apply, is a storable fact. Storable facts are grammar-relative: they are constituted within the identity criteria, distinctions, and lawful relationships of some grammar, and their probability of accurate correspondence is conditioned on that grammar. Therefore JTB's truth condition, which requires grammar-independent truth, requires a property that no storable fact possesses. This is revisionary enough to defeat JTB as an analysis of knowledge without committing to strong constructivism. The argument does not need the strong claim. The weak claim, properly followed through, is sufficient.

This definition also unifies what might otherwise appear to be disparate uses: the geocentric grammar, the RCT grammar, the grammar of classical mechanics. All are instances of the same structure. All generate remainder. All are subject to supersession. The term carries the same meaning in each case.

We develop the argument in six stages. Section 2 identifies three structural features of conscious knowing that characterize the epistemic condition JTB cannot describe. Section 3 demonstrates the structure through the Copernican revolution and engages Kuhn's account of paradigm change directly. Section 4 shows the structure operating in compound clinical knowledge claims, where inferential absence renders claims formally unevaluable. Section 5 explains why existing frameworks, including Bayesian epistemology and imprecise probability approaches, cannot name the condition we identify. Section 6 proposes Justified Probable Belief as the replacement framework. Section 7 concludes.

2. Three Structural Features of Conscious Knowing

The epistemic condition of conscious knowers is shaped by three structural features that JTB has no vocabulary for. Each is independently established. Together they characterize what knowing actually is for beings like us, and together they explain why JTB's truth condition cannot serve as a description of the epistemic achievement conscious knowers make.

2.1 The Grammar-Entry of Belief Formation

Knowledge is formed within conscious experience. Conscious experience has a temporal structure with a precise consequence for any account of knowledge: the moment at which a knower has the most direct access to a state of affairs is the moment of immediate experience, and that moment does not hold still.

By the time a belief is formed about an experience, the experience is already becoming memory. Memory is not a faithful recording. It is an active reconstruction shaped by subsequent experience, expectation, and the conceptual frameworks the knower brings to their own history. The belief is about a reconstruction of a state of affairs, not the state of affairs itself.

This observation does not require that knowledge of recently past events is impossible, or that JTB must freeze a vanishing present. Propositions can be about what just happened. The point is narrower and more precise: the act of forming a belief is already an act of grammatical constitution. The knower cannot form a belief about an experience without organizing it, and organization is always grammar-relative. The identity criteria, distinctions, and lawful relationships of the grammar the knower inhabits are the tools with which experience is transformed into a storable belief. A knower who inhabits a different grammar would constitute the same experience into a different belief. The experience does not determine the belief independently of the grammar. The grammar enters the belief at formation.

We call this the grammar-entry of belief formation. It is not a claim about fallibility in the ordinary sense. It is a claim that every belief is grammatically constituted from its first moment, and that its correspondence to any state of affairs is correspondence within the grammar, not across all possible grammars. JTB's truth condition applies to the belief as if it were a report of a grammar-independent state of affairs. It is not. It is a grammatically constituted representation, and the grammar-independence that JTB's truth condition requires is not available at the point of formation.

2.2 The Directional Drift of Linguistic Transmission

Every knowledge claim of any complexity is communicated through language. Language has a structural property that JTB has no framework for: the transmission of a belief through language is not a transfer of content but a transformation of it, and the transformation is systematic and directional.

Meaning operates along two dimensions that stand in permanent tension. The first is definitional: the abstract, context-independent content of a term that allows a listener who does not share the speaker's situation to understand it. The second is contextual: the particular, situation-specific force a term carries in this instance of use. No formulation can fully achieve both simultaneously. A commitment to definitional precision severs the term from the context that gave it specific force. A commitment to contextual fidelity resists the fixed definition that would allow accurate decoding.

The direction of drift under ordinary communicative pressure is not random, not merely anecdotal, and not confined to any single domain. It has a structural basis in the architecture of meaning itself. Grice's cooperative principle, articulated in *Logic and Conversation* (1975) and developed in *Studies in the Way of Words* (1989), identifies a set of maxims that govern rational communicative exchange. The maxim of quantity requires speakers to make their contribution as informative as is required for the current purposes of the exchange. The maxim of manner requires clarity and the avoidance of obscurity. In action-oriented communicative contexts, these two maxims apply systematic pressure toward the most informative, least hedged reading the evidence can support. Probabilistic qualifiers, methodological caveats, and scope conditions reduce informativeness and introduce the appearance of obscurity. They are therefore systematically deprioritized under communicative pressure, not through carelessness but through the rational structure of meaning exchange itself. This is not a domain-specific finding. It is a consequence of the architecture of cooperative communication as Grice describes it, operating across all communicative contexts where the purpose is action rather than epistemic accuracy.

The consequence of this structural pressure has been documented empirically in the domain where it carries the highest stakes. Woloshin and Schwartz's systematic research on health communication showed that medical findings are consistently amplified as they travel from trial results through press releases to news coverage, with each stage introducing stronger claims, dropping qualifications, and omitting hedges that were present at the source (Woloshin and Schwartz, 2006; Schwartz and Woloshin, 2002). The pattern is directional: toward stronger readings, higher confidence, more actionable conclusions. Associated with becomes causes. May reduce becomes reduces. Statistically significant becomes important. Grice explains why this pattern is structural rather than contingent: under the communicative pressure of action-oriented exchange, any formulation that preserves the hedges is less informative by the maxim of quantity and less clear by the maxim of manner than a formulation that drops them. The drift is not a failure of the system. It is the system working as designed, for purposes other than epistemic fidelity.

The consequence for JTB is not that perfect transmission is logically impossible. It is that JTB has no account of what happens to the truth condition during transmission. JTB applies its conditions to a belief. But the belief that arrives at a listener after transmission through language is not the belief that was sent. The grammar-entry that constituted the belief at formation is compounded by the directional drift that transforms it during transmission. JTB treats the transmitted claim as epistemically equivalent to the original. It is not. An adequate epistemological framework must represent the gap between them. JTB cannot.

2.3 The Grammar-Relativity of Facts

The third structural feature is the most fundamental and applies to the truth condition most directly. Every knowledge claim is constituted within a grammar as defined in section 1. This has a precise consequence for JTB's truth condition.

JTB requires that if P is true, it is true regardless of the framework within which P is articulated. Truth, on the JTB account, is grammar-independent. But every storable fact is constituted within a grammar. Its probability of accurate correspondence is the probability of its accuracy within the current closure regime, conditioned on that regime's identity criteria, distinctions, and lawful relationships. A claim that achieves maximum probability within one grammar can become a different kind of claim entirely within the next. JTB's truth condition is binary and grammar-independent. The actual epistemic condition of knowledge claims is probabilistic and grammar-relative. JTB cannot distinguish between these two conditions because it has no concept of grammar in the sense we have defined it.

This is not a claim that truth does not exist as a property of propositions. It is a claim that JTB's truth condition cannot serve as a description of what conscious knowers have when they know, because the condition they occupy is grammar-relative in a way that the binary truth condition cannot accommodate or even see.

3. The Copernican Demonstration and the Extension of Kuhn

The three structural features identified above are philosophical arguments about the epistemic condition of grammar-inhabiting knowers. The most powerful demonstration of their consequences is historical. We enter this demonstration with the framing from section 1 clearly in

view: what we are showing is not that JTB fails to be a theory of scientific progress, but that JTB cannot describe what the Ptolemaic astronomer had epistemically, which is a failure of JTB as an analysis of knowledge for conscious knowers. The Copernican case is a demonstration of descriptive inadequacy, not a demand that JTB serve as a philosophy of science.

By the second century CE, Claudius Ptolemy had developed a mathematical model of planetary motion that was the most precisely verified scientific framework of the ancient and medieval world. Ptolemaic astronomy predicted celestial positions with accuracy sufficient for navigation, agriculture, and the calculation of eclipses. It had been tested against observation for over a thousand years. Within the geocentric grammar, the probability of the core claim approached certainty. By any reasonable epistemic standard, the Ptolemaic astronomers knew what they claimed to know: their belief was justified by extensive replicated predictively successful evidence, sincerely held, and by the realist's own account about a real world in which planets moved and eclipses occurred.

JTB's verdict on this case is: the geocentric belief was false, because the earth orbits the sun. That verdict is available to JTB and it is not wrong. But it does no explanatory work. It does not explain what the Ptolemaic astronomers had, why the revision happened, what it means that their predictions worked, or how science progresses through episodes in which maximally justified beliefs are superseded without error being discovered. A framework that can only describe the most consequential episode in the history of science as: they were wrong, is not an adequate framework for understanding knowledge.

The Copernican revision did not work the way JTB imagines knowledge being corrected. The Ptolemaic astronomers had not made a factual error within their grammar. Their observations were accurate. Their mathematics was valid. Their predictions worked. The sun still rises and sets. The observations that justified the geocentric model are the same observations that justify the heliocentric model. What changed was not the data. What changed was the grammar: the identity criteria that determined what counted as the center of the system, the distinctions that made certain relationships legible, the lawful relationships that organized observations into predictions. Copernicus reorganized existing observations within a new closure regime. The remainder the Ptolemaic grammar had accumulated dissolved. The new grammar was more adequate to the world it was modeling.

This is the point at which Kuhn's account is both our closest predecessor and the place where our framework makes a distinct contribution.

Kuhn identified that scientific revolutions involve paradigm change rather than simple accumulation of evidence, that observation is theory-laden, and that successive paradigms may be incommensurable in ways that make rational comparison difficult (Kuhn, 1962). These are genuine insights that our framework inherits and builds on. The geocentric and heliocentric systems are exactly the kind of case Kuhn analyzed. But the concept of remainder we introduce is not simply a relabeling of the anomalies Kuhn described. The distinction matters and must be stated precisely.

Anomalies, in Kuhn's account, are observations that resist a paradigm's explanatory resources: they are contingent features of what a particular paradigm happens to encounter. Remainder, as defined in this paper, is a structural property of the grammar itself: it is the mismatch between what any finite closure regime can model and what the world contains, generated necessarily by the grammar's own finite identity criteria. Anomalies are symptoms. Remainder is

the underlying condition that makes anomalies possible. A grammar that encountered no anomalous observations would still generate remainder, because the structural mismatch between finite identity criteria and an inexhaustible world is not contingent on any particular observation. This distinction transforms Kuhn's historical account into a structural prediction: not just that paradigms face anomalies, but that every grammar necessarily generates remainder and therefore necessarily faces supersession pressure, regardless of what observations happen to be made.

Our framework extends Kuhn's account in two specific ways. First, Kuhn identified the phenomenon of paradigm change intuitively without providing a generative mechanism. Our concept of remainder provides that mechanism structurally rather than historically. Second, Kuhn's account left the epistemological consequence of paradigm change underspecified. He showed that science changes in ways that look irrational by the old paradigm's standards, but he did not provide a replacement account of what scientists have when they have knowledge within any paradigm. Justified Probable Belief provides that account. Kuhn is the diagnosis. JPB is the framework the diagnosis requires.

Our framework also addresses the incommensurability challenge more precisely than Kuhn's original account. Grammars can be compared by the adequacy of their absorption of remainder: a new grammar that dissolves the remainder the old grammar accumulated is demonstrably more adequate to the world, even if the comparison cannot be made from within either grammar alone. This gives rational theory choice a criterion that survives incommensurability without requiring a grammar-independent vantage point. The criterion is the best available grammar: the one that has dissolved the most remainder relative to its predecessors. This is not circular. It is the same criterion by which scientists in practice evaluate competing frameworks, now given a structural name.

The heliocentric grammar was itself superseded. Kepler replaced circular orbits with ellipses. Newton replaced geometric descriptions with gravitational mechanics. Einstein replaced absolute space with relativistic spacetime. Each revision dissolved remainder the previous grammar had accumulated. Each produced a grammar more adequate to the world. None of this required access to grammar-independent truth. It required the accumulation and dissolution of remainder. That is what scientific progress is. JTB has no account of it. JPB describes it exactly.

What the Copernican case demonstrates is not that the Ptolemaic astronomers made an error. It is that JTB cannot distinguish between what they had and what a modern physicist has. Both satisfy JTB's conditions when their beliefs are maximally justified and sincerely held. JTB treats them as epistemically equivalent. They are not. One is operating within a grammar that has been superseded. The other is operating within a grammar that is the current best. That difference is what matters for understanding knowledge, and JTB cannot see it.

4. The Clinical Demonstration: Inferential Absence

The three structural features of conscious knowing apply to all knowledge claims made by beings with the properties conscious knowers have. They are most visible in domains where inferential chains are long, stakes are high, and the gap between what evidence establishes and what claims assert is both structurally precise and largely invisible to existing frameworks. Clinical medicine is that domain. What follows is not an empirical proof of the philosophical case, which was made in sections 1 through 3. It is a demonstration of inferential absence in the domain where

it carries the most practical consequence for human life. The philosophical argument stands independently of whether any particular empirical detail in this section is contested.

4.1 Compound Claims and Inferential Absence

A compound clinical claim of the form this intervention benefits patients with this condition decomposes into a chain of inferential steps, each of which must hold for the compound to hold, and each of which has an epistemic state that can be examined independently.

Each inferential step falls into one of three categories. Grounded: the step has an assignable probability derived from replicated evidence with known scope conditions. Estimated: the step has a range derived from partial or contested evidence, acknowledged as uncertain. Structurally unmeasurable: no study exists from which any probability or probability bound could be derived, either because the relevant mapping has not been studied or because it cannot be studied within the current methodological grammar.

The formal argument for why structurally unmeasurable components terminate the compound evaluation rather than merely widening the uncertainty interval requires two steps. First: a structurally unmeasurable component has no evidential basis for any boundary on its probability. It is not a wide credal set. It is the absence of any constraint that would define a set at all. Any number assigned to it is not derived from evidence but generated without inferential basis. It is an invented number, not an estimated one. Second: multiplying genuine probability estimates by invented numbers does not produce a probability estimate. It produces a number that carries the form of a probability and none of its epistemic substance. The product is not a wide interval or a poorly grounded estimate. It is undefined: a term that describes not extreme uncertainty but the absence of the inferential structure that probability estimates require. This is not a strong claim about the world. It is a precise description of what the arithmetic contains.

The compound structure analysis that follows is presented as a diagnostic framework, not a precision quantitative model. Where we assign numerical probabilities to grounded components, those numbers are illustrative estimates consistent with the published literature, not precise calculations. The diagnostic power of the framework lies in identifying which components are grounded, which are estimated with acknowledged uncertainty, and which are structurally unmeasurable. The arithmetic is secondary to that categorization. What the multiplication demonstrates is the qualitative structure of the compound claim: how many links the chain contains, which links have any evidential basis, and why the presence of structurally unmeasurable links renders the compound unevaluable regardless of the precise values assigned to the grounded ones.

This condition, which we call inferential absence, must be distinguished from the cases handled by imprecise probability frameworks and by transportability and partial identification analysis. We address those frameworks directly in section 5.

4.2 The Procedure for Compound Claim Analysis

Before applying the framework to a specific claim, the procedure requires three steps.

First, identify every inferential step required for the conclusion to follow from the evidence. A compound claim is not a single assertion. It is a chain, and every link must be named explicitly. Each link is a separable proposition with its own epistemic state.

Second, type each component as grounded, estimated, or structurally unmeasurable, with explicit statement of the basis for each typing. Grounded components require citation of the replicated evidence and its scope conditions. Estimated components require acknowledgment of the uncertainty and its sources. Structurally unmeasurable components require identification of what study would be required to make them measurable and why that study has not been or cannot be conducted within the current methodological grammar.

Third, apply the multiplication. The product of grounded and estimated components is the maximum calculable confidence level the compound can honestly claim. Any structurally unmeasurable component terminates the multiplication. The compound probability is undefined. Assigning it a confidence level anyway is not probabilistic reasoning under uncertainty. It is what we call the confidence surplus: a number generated without inferential basis, formatted as if it had one.

4.3 A Worked Demonstration: The Grammar of Depression Treatment

The following decomposition applies the compound structure to the regulatory claim that antidepressant medication treats depression. This case is chosen because it is the domain where the RCT grammar's remainder is most visible, most documented, and most consequential for everyday life. The researchers who conducted these trials were rigorous. The regulators who evaluated them applied the best available standards. The clinicians who prescribe these medications are trying to help their patients. The patients who take them often do improve. None of that is in dispute. What is in dispute is whether the grammar of the RCT can certify the compound claim the approval is taken to establish.

Before the decomposition, one preliminary observation establishes why depression is the terminal demonstration of grammar generating remainder from the inside rather than at its edges. Antidepressant drug arm response rates and placebo arm response rates track each other so closely across trials, drug classes, severity levels, and treatment modalities that the gap between them consistently falls below the threshold of clinical significance. This pattern holds across individual trials, across meta-analyses of published data, and across the unpublished FDA submission data analyzed by Kirsch and colleagues (Kirsch et al., 2008). The grammar of the RCT predicts that pharmacological action should be additive to expectation-driven response and should produce consistent and substantial separation between arms. It does not. We return to this finding after the decomposition, because it is the remainder made undeniable.

The compound claim: This molecule treats depression.

C1: The trial produced a statistically significant difference between the drug arm and the placebo arm on the specified endpoint.

Type: Grounded. This is what the trial actually establishes and establishes well. The statistics are internally valid. The significance threshold was met. We assign this a probability of approximately 0.85, reflecting genuine statistical confidence discounted for

publication bias and replication variance across trials. The grammar works here. The problem begins with every step that follows.

C2 | C1: The blinding was maintained with sufficient integrity that the drug arm result reflects pharmacological action rather than expectation-driven response.

Type: Structurally unmeasurable within the current trial design. Depression drugs produce side effects. Patients on the drug arm experience those side effects and correctly infer with high probability that they are receiving active treatment. That inference activates auto-noetic engagement: the patient holds a representation of themselves as someone receiving a medication that is producing detectable physiological effects. That representation is a biological event activating expectation-driven recovery mechanisms that the placebo arm was designed to hold constant. Research by Rabkin, Margraf, and subsequent investigators has consistently shown that patients in antidepressant trials correctly identify their arm assignment at rates substantially above chance. A 2022 systematic review found that only 4.7 percent of antidepressant RCTs formally assessed blinding integrity, while a second 2022 review found that in newer-antidepressant trials patients and assessors were unlikely to correctly judge allocation, presenting a mixed record on the magnitude of unblinding effects. Active-placebo and placebo run-in designs represent genuine attempts within the RCT grammar to address the blinding problem. The argument here is not that such attempts do not exist, but that they have not been shown to resolve the structural question of how much of the drug arm result reflects pharmacological action versus expectation-driven response in standard populations under standard protocols. The degree to which expectation inflation contributed to the drug arm result in any given standard trial cannot be measured by that trial's design, because the design has no arm that received pharmacological action without the expectation that typically accompanies active treatment. That specific decomposition remains structurally unmeasurable within the standard trial grammar. $P(C2 | C1) =$ structurally unmeasurable within standard trial designs.

C3 | C1, C2: The placebo arm endpoint rate reflects the absence of pharmacological treatment rather than the presence of nocebo suppression.

Type: Structurally unmeasurable within the current trial design. Patients on the placebo arm who suspect they are receiving inactive treatment may actively underperform relative to their biological potential because reduced expectation suppresses the recovery mechanisms both arms depend on. The brain representing itself as untreated is not neurologically neutral. Nocebo suppression is a real and documented phenomenon in other clinical contexts. Its magnitude within any specific antidepressant trial is not measured because the trial has no mechanism to separate spontaneous recovery, regression to the mean, and nocebo-suppressed recovery within the placebo arm result. The gap between arms is not a clean measure of pharmacological action. It is the difference between a potentially expectation-inflated drug arm and a potentially nocebo-suppressed placebo arm. Both biases favor the drug arm. Neither is measurable within the current trial grammar. $P(C3 | C1, C2) =$ structurally unmeasurable.

C4 | C1, C2, C3: The observed endpoint gap reflects a pharmacological mechanism that will operate in the prescribing context as it operated in the trial.

Type: Structurally unmeasurable. The trial held clinical context as constant as possible. It did so deliberately: auto-noetic amplification is the variable the trial was designed to suppress so that the pharmacological contribution could be measured against a controlled baseline. The prescribing context is not merely different. It is the systematic opposite. Phase 3 marketing, drug naming, diagnostic narrative, clinical training, and the ongoing therapeutic relationship are all instruments of auto-noetic amplification, deployed after approval to maximize the very mechanisms the trial worked to flatten. The trial grammar suppresses auto-noetic engagement in order to measure. The commercial grammar maximizes it in order to deploy. These are opposing operations applied to the same mechanism. The regulatory certification produced under suppressed conditions is then used to authorize prescribing under maximally amplified conditions. The grammar treats them as equivalent contexts because both fall under the same approval. The biology recognizes no such equivalence. The conditions are not merely different but structurally inverted. $P(C4 | C1, C2, C3) = \text{structurally unmeasurable}$.

C5 | C1...C4: The pattern of drug and placebo arm convergence across trials is consistent with the claim that the molecule produces pharmacological action independent of and additive to expectation-driven mechanisms.

Type: Structurally unmeasurable. This is the terminal demonstration. If pharmacological action were the primary driver of drug arm response and genuinely additive to expectation-driven recovery, drug response rates should produce consistent and substantial separation from placebo arm results. They do not. Drug and placebo arms track each other so closely across trials, drug classes, severity levels, and treatment modalities that the gap consistently falls below the threshold of clinical significance. The grammar predicts separation. The data refuses it. Kirsch and colleagues' analyses of FDA submission data, as well as subsequent meta-analyses, document this convergence across multiple independent datasets (Kirsch et al., 2008; Cipriani et al., 2018). The counter-argument is that expectation is itself a treatment mechanism and activating it is a legitimate clinical achievement. That argument may be correct. It cannot be resolved within the current grammar because the grammar cannot decompose its own primary finding into pharmacological and expectation-driven constituents. That irresolvability is not a gap awaiting more data. It is remainder: the structural mismatch between what the RCT grammar can measure and what the biology of depression treatment requires. $P(C5) = \text{structurally unmeasurable}$.

Applying the compound structure:

$$\begin{aligned} P(K) &= P(C1) \times P(C2|C1) \times P(C3|C1,C2) \times P(C4|C1,C2,C3) \times \\ &P(C5|\dots) \\ P(K) &= 0.85 \times \text{unmeasurable} \times \text{unmeasurable} \times \text{unmeasurable} \times \\ &\text{unmeasurable} \\ &= \text{undefined} \end{aligned}$$

The compound is formally unevaluable from the first structurally unmeasurable component forward. The only grounded component, the internal validity of the trial result, sits at 0.85. Every subsequent inferential step is structurally unmeasurable within the current grammar. The compound probability of the claim this molecule treats depression is formally unevaluable not because the science is weak but because the grammar of the RCT cannot measure what the clinical claim requires.

The confidence at which this claim is presented to regulators, to prescribers, to patients approaches certainty. The maximum calculable probability of its foundational component is 0.85. Every component connecting that finding to the clinical claim is structurally unmeasurable. The confidence surplus, the quantity of unwarranted confidence that the existing epistemic framework cannot identify or name, is the distance between those two positions.

Now return to the convergence pattern. Kirsch and colleagues have argued that the net pharmacological contribution of approved antidepressants may be clinically modest (Kirsch et al., 2008). The counter-argument within the existing grammar is that expectation is itself a treatment mechanism. Both arguments may be correct. Neither can be resolved within the RCT grammar, because the grammar cannot decompose its own primary finding into its constituents. That irresolvability is remainder. The convergence of drug and placebo arms across trials, classes, and severity levels is the world exceeding the grammar's capacity to describe it. The grammar has generated remainder at its center, not at its edges. The remainder has accumulated sufficient force that it has a name: the placebo problem. What it has not yet generated is a new grammar adequate to replace the one it is straining against.

4.4 The Structure Is Not Unique to Clinical Evidence

The condition of inferential absence is most visible in clinical evidence because the inferential chains are long, the stakes are high, and the absent components are consequential enough to have generated a substantial critical literature without being named. But the structure is not specific to medicine. It is the compound expression of the same grammar-relativity and remainder-generation that characterize all knowledge claims made within finite closure regimes.

Any compound knowledge claim that traverses inferential steps connecting evidence to conclusion can contain structurally unmeasurable components. Policy recommendations derived from economic models, historical causal claims, projections in climate science, behavioral predictions in social science: all involve compound inferential chains that may contain steps for which no probability or probability bound is available, not because the evidence is weak but because the relevant study has not been undertaken or cannot be undertaken within the current grammar. JTB has no framework for any of these cases. The condition of inferential absence does not exist within JTB's vocabulary.

5. Why Existing Frameworks Cannot Name This

5.1 The Clinical Evidence Literature

Cartwright, Howick, Worrall, and the broader literature on evidence in clinical epistemology have produced the most sophisticated available critique of how RCT findings travel

from experimental populations to clinical practice. Each identifies a genuine problem. None names inferential absence, because each operates within a framework that assumes the compound claim has an evaluable probability, however uncertain or poorly grounded.

Cartwright's argument that RCTs establish that an intervention worked somewhere under some conditions, not that it will work in a given target population, is a direct ancestor of the compound structure analysis this paper proposes (Cartwright, 2007). Her framework identifies the generalizability gap: the inferential step from trial to target population is epistemically non-trivial and the hierarchy's ranking of RCTs systematically underweights its difficulty. This paper builds on that insight and shows precisely where it stops short. Cartwright's framework, even fully applied, addresses cases of imperfect generalization, where some structure connects trial to target but the connection is incomplete. It has no resources for the prior condition in which the connection has not been established at all and cannot be established within the current grammar.

Worrall's mechanistic argument holds that without knowing why an intervention works, evidential support for clinical application is weaker than the hierarchy implies (Worrall, 2007). This is correct. The connection to this paper's argument is that the mechanistic question presupposes the trial result represents genuine causal action rather than expectation-driven effects. In trials where blinding is compromised by side effects, that presupposition is itself structurally unmeasurable. Worrall identifies a problem at a level of the inferential chain that presupposes a prior step has been resolved. That prior step is itself structurally unmeasurable.

Howick's challenge to the hierarchy argues that mechanistic reasoning is systematically undervalued relative to RCTs (Howick, 2011). This paper takes no position on the relative ranking of evidence types. Its argument is that neither RCTs nor mechanistic reasoning, at any level of the hierarchy, can assign a probability or probability bound to inferential steps that have not been studied. The hierarchy question is downstream of the evaluability question.

5.2 Bayesian Epistemology and Imprecise Probability

The natural formal response to our argument is Bayesian. Bayesian agents update credences on evidence, handle uncertainty through probability assignments, and accommodate conditional probabilities exactly as the compound structure requires. And the sophisticated Bayesian response to our critique is imprecise probability: credal sets, interval-valued beliefs, and suspension of judgment frameworks developed by Walley and extended by subsequent researchers (Walley, 1991). These frameworks explicitly distinguish evidence that supports a credence of 0.5 from evidence that is silent between 0 and 1. They are designed precisely to avoid pretending to know more than the data support while still saying something disciplined.

We engage this literature directly because it is the strongest available response to our argument, and because the distinction between what it handles and what we identify is philosophically important.

Imprecise probability frameworks handle cases where the evidence is insufficient to warrant a single sharp credence but is sufficient to bound the interval. A credal set is still a set: it has boundaries, however wide. The framework requires some evidential basis for those boundaries, even if the basis is weak or contested. What it cannot handle is the condition where no basis exists for any boundary at all: not a wide interval, but the absence of any structure that would constrain the interval. That is the condition inferential absence names. When an inferential step has not been

studied and cannot be studied within the current methodological grammar, there is no evidential basis for even the widest credal set. The interval is not wide. It is empty. Assigning any interval in that condition is not imprecise reasoning. It is the invention of a structure that does not exist, formatted as if it did.

The distinction matters because the two conditions require different responses. Imprecise probability is the right framework for C1 in the depression decomposition: the trial result has grounded evidential basis, the uncertainty about publication bias and replication variance can be bounded, and a credal set is the appropriate representation. It is not the right framework for C2 through C5, because those steps have no structure connecting trial conditions to the relevant target. No selection diagram has been drawn. No constraints on the identified set have been established. The mapping has not begun. Imprecise probability frameworks presuppose that mapping has begun, however imperfectly. Inferential absence is the condition where it has not.

Transportability analysis and partial identification frameworks address the problem of generalizing from trial populations to target populations under conditions of imperfect mapping (Pearl and Bareinboim, 2014). These are powerful tools for the cases Cartwright identifies. But they require that some structure connecting trial to target has been established, even if incompletely. Transportability analysis requires selection diagrams specifying which variables differ between trial and target. Partial identification requires at least some constraints on the identified set. Inferential absence is the condition where that structure has not been established at all. These frameworks cannot operate where their required inputs do not exist.

JTB has the same structural limitation at a more fundamental level. It presupposes that the belief either corresponds to the state of affairs or does not, a binary that assumes the correspondence is assessable in principle. When the correspondence is grammar-relative and the grammar is finite and revisable, JTB's binary produces a verdict where the honest answer is: this claim cannot be evaluated within the current grammar. That is inferential absence at the level of all knowledge, of which the clinical case is the most visible instance.

6. Justified Probable Belief

If JTB misdescribes what conscious knowers actually have, what is the accurate description? We propose Justified Probable Belief as a framework proposal, not a completed formal theory. We present it with explicit acknowledgment of what a completion would require, because intellectual honesty about the framework's current state is itself an application of the epistemic standard we are proposing.

A self-application objection requires a direct answer before proceeding. If all knowledge claims are grammar-relative, is JPB itself only valid relative to a grammar? If yes, then its critique of JTB has only local force: JPB would be the better framework within our current grammar but could not claim universal superiority. If no, then at least one claim escapes grammar-relativity, which appears to undercut the paper's universal thesis.

The answer is that JPB is a second-order framework about the structure of knowing, not a first-order knowledge claim about the world. Its grammar-relativity thesis is structural and self-consistent: it claims that all first-order knowledge claims are grammar-relative, which is itself a structural claim about the form of knowledge rather than a claim about any particular domain of

facts. This is analogous to how a theory of logic can state logical laws without itself needing to stand outside logic. A claim about the form that all knowledge takes does not require exempting itself from that form. JPB's own status is: maximum probable belief within the current best grammar for epistemological theorizing, held honestly, open to revision when a better grammar of epistemology becomes available. That is not a weakness. It is the framework applying itself consistently. JTB, by contrast, presents its own conditions as grammar-independent, which is precisely the move JPB identifies as unavailable to any grammar-inhabiting knower.

A person S has justified probable belief that P to degree $P(K|G)$ if and only if S believes P, S is justified in believing P, and $P(K|G)$ is the probability that P accurately represents the relevant state of affairs, conditioned on the current grammar G and explicitly open to revision when G is superseded.

The probability type of $P(K|G)$ requires explicit specification because the objection that it is unclear whether this is subjective credence, evidential probability, objective chance, or intersubjective warrant is a legitimate one. $P(K|G)$ is the intersubjective degree of warranted acceptance of P within the community of competent inquirers operating within grammar G, assessed against the standards of evidence, inference, and coherence that grammar G makes available. It is not subjective credence, which varies by individual and does not require community assessment. It is not objective chance, which would require grammar-independence. It is not formal evidential probability in the Bayesian sense, which requires a prior and does not condition on a grammar in the sense we have defined. It is the community-level epistemic standing of P within G: the best answer available to the question of how strongly the grammar warrants this claim, assessed by competent inquirers who share the grammar's identity criteria and standards of evidence. This interpretation is coherent, non-circular, and consistent with how scientific communities actually assess knowledge claims. It also explains why JPB preserves the normative hierarchy: within any grammar, some claims have higher $P(K|G)$ than others, and that ranking is the epistemic work that truth was supposed to do but, for reasons the paper has argued, cannot do grammar-independently.

Grammar G is individuated by the definition given in section 1: a stabilized closure regime identified by its identity criteria, distinctions, and lawful relationships. Two descriptions belong to the same grammar when they share those constitutive elements. A grammar is superseded when a new closure regime dissolves the remainder the old grammar had accumulated. The best available grammar is the one that has dissolved the most remainder relative to its predecessors. This criterion is not circular: it is assessed by examining what anomalies and irresolvable discrepancies the old grammar accumulated and whether the new grammar has resolved them. $P(K|G)$ is a probability estimate about the world as constituted by grammar G, with the understanding that G is incomplete and that a new grammar is always possible.

Four features of this formulation require attention.

First, the probability is conditioned on the grammar. This is not relativism. Within any given grammar, claims can be ranked by probability, and higher probability is epistemically superior. The heliocentric model has a higher $P(K|G)$ within the current grammar of physics than any geocentric model. JPB preserves the normative hierarchy JTB was designed to capture. What it abandons is the description of that hierarchy as terminating in grammar-independent truth rather than in maximum probability within the best available grammar. That is not a weakening of epistemic standards. It is an accurate description of what the hierarchy actually delivers.

Second, JPB introduces a formal category JTB lacks: the unevaluable claim. Under JTB, claims are true or false. Under JPB, claims occupy positions on a probability spectrum conditioned on the current grammar, and some compound claims cannot occupy any position because their inferential chains contain structurally unmeasurable components. These claims are not false. They are not uncertain. They are unevaluable, and JPB requires that they be identified and reported as such.

Third, the update rule under grammar change is: $P(K|G)$ does not constrain $P(K|G')$ when G is superseded by G' . A high-probability claim in one grammar can become a low-probability claim in the next. This is not a deficiency. It is the honest acknowledgment that grammar supersession is a genuine epistemic event, not a correction of error. What carries over across grammar change is the evidence: the observations remain. What changes is the grammar that constitutes those observations into facts.

Fourth, JPB does not yet provide a complete formal theory. Three open questions require resolution by any adequate completion: a precise decision procedure for identifying grammar boundaries, a formal account of when remainder crosses the threshold that makes supersession possible, and a rigorous treatment of how $P(K|G)$ values are aggregated across knowers within a grammar. We identify these as open not to disclaim the framework but to specify what the research program requires.

The same analysis applies to falsity. Under JTB, false is the complement of true. But grammar revision makes falsity as grammar-relative as truth. Within the Ptolemaic grammar, the earth moves around the sun was assigned probability approaching zero. Grammar revision made it the higher-probability claim. Near-zero within the current grammar is not grammar-independent falsity. Both poles of JTB's binary are grammar-relative approximations. Neither is reachable with certainty. JPB names that condition honestly rather than pretending the binary is available.

JPB is not a diminishment of epistemic achievement. Science in its rigorous practice has always operated as a JPB discipline. P-values, confidence intervals, replication requirements, and the provisional status of all theories are institutional expressions of a JPB epistemology that has never been given its proper philosophical name. Closing the gap between what science produces and what JTB claims it produces is not a retreat from rigor. It is what rigor actually requires.

7. Conclusion

Justified True Belief correctly identifies truth as the condition that would distinguish knowledge from mere belief. What it cannot do is provide any account of how beings who form beliefs within transitory experience, transmit them through language, and operate within finite revisable grammars could distinguish truth from the condition they actually occupy: maximum probable belief within the best available grammar.

The grammar-entry of belief formation establishes that every knowledge claim is grammatically constituted from its first moment. The directional drift of linguistic transmission establishes that the grammar-relativity introduced at formation is compounded and amplified during communication, in a direction documented by systematic research on how medical evidence travels from trials to public claims. The grammar-relativity of facts establishes that every claim's probability of accurate correspondence is conditioned on the closure regime that constitutes it as a

statable fact. JTB's binary truth condition has no vocabulary for any of these features of the epistemic condition of conscious knowers.

The Ptolemaic astronomers satisfied every condition JTB specifies. JTB's only available description is that their belief was false. That description does no explanatory work. Our framework explains what JTB cannot: they had maximum probable belief within their grammar, the revision happened because remainder accumulated to the point where a new grammar became possible, and scientific progress is the successive supersession of grammars by ones more adequate to the world they model. Kuhn described this process. Our framework formalizes the mechanism and supplies the epistemological consequence his account left open.

In clinical medicine the same structure operates where it matters most. Compound inferential claims connecting trial findings to treatment decisions contain structurally unmeasurable components: inferential steps for which no probability or probability bound can be derived within the current methodological grammar. The confidence surplus between how such claims present themselves and what their evaluable components can support is large, consequential, and invisible to every existing framework, including Bayesian epistemology and imprecise probability approaches, because none of those frameworks has a formal category for a claim that is not uncertain but unevaluable.

Justified Probable Belief describes what conscious knowers actually have. Positions on a probability spectrum, conditioned on the best available grammar, held honestly, open to revision when the grammar is superseded, and disciplined enough to say when a claim cannot be evaluated at all. Near-zero is what we have been calling false. Near-one is what we have been calling true. Neither end is reachable with certainty. Both are grammar-relative. Both are open to revision.

JTB asked for something the structure of conscious knowing cannot deliver. JPB describes what conscious knowing actually delivers. That is not a concession. It is what intellectual honesty, rigorously applied, looks like.

References

- Cartwright, N. (2007). *Hunting Causes and Using Them: Approaches in Philosophy and Economics*. Cambridge University Press.
- Cipriani, A., Furukawa, T. A., Salanti, G., et al. (2018). Comparative efficacy and acceptability of 21 antidepressant drugs for the acute treatment of adults with major depressive disorder: a systematic review and network meta-analysis. *The Lancet*, 391(10128), 1357-1366.
- Gettier, E. L. (1963). Is Justified True Belief Knowledge? *Analysis*, 23(6), 121-123.
- Goldman, A. I. (1979). What Is Justified Belief? In G. Pappas (Ed.), *Justification and Knowledge*. Reidel.
- Grice, H. P. (1975). Logic and Conversation. In P. Cole and J. Morgan (Eds.), *Syntax and Semantics*, Vol. 3: *Speech Acts* (pp. 41-58). Academic Press.

- Grice, H. P. (1989). *Studies in the Way of Words*. Harvard University Press.
- Howick, J. (2011). *The Philosophy of Evidence-Based Medicine*. Wiley-Blackwell.
- Kirsch, I., Deacon, B. J., Huedo-Medina, T. B., Scoboria, A., Moore, T. J., and Johnson, B. T. (2008). Initial Severity and Antidepressant Benefits: A Meta-Analysis of Data Submitted to the Food and Drug Administration. *PLoS Medicine*, 5(2), e45.
- Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.
- Margraf, J., Ehlers, A., Roth, W. T., et al. (1991). How 'blind' are double-blind studies? *Journal of Consulting and Clinical Psychology*, 59(1), 184-187.
- National Institute for Health and Care Excellence (NICE). (2022). *Depression in Adults: Treatment and Management*. NICE Guideline NG222. NICE.
- Pearl, J., and Bareinboim, E. (2014). External Validity: From Do-Calculus to Transportability Across Populations. *Statistical Science*, 29(4), 579-595.
- Rabkin, J. G., Markowitz, J. S., Stewart, J., et al. (1986). How blind is blind? Assessment of patient and doctor medication guesses in a placebo-controlled trial of imipramine and phenelzine. *Psychiatry Research*, 19(1), 75-86.
- Schwartz, L. M., and Woloshin, S. (2002). Media coverage of scientific meetings: too much, too soon? *JAMA*, 287(21), 2859-2863.
- Sild, A. I., Haile, C. N., and Grabowski, J. (2022). Assessment of blinding in antidepressant randomized controlled trials: a systematic review. *Journal of Clinical Psychiatry*, 83(5).
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press.
- Woloshin, S., and Schwartz, L. M. (2006). Media reporting on research presented at scientific meetings: more caution needed. *Medical Journal of Australia*, 184(11), 576-580.
- Worrall, J. (2007). *Evidence in Medicine and Evidence-Based Medicine*. *Philosophy Compass*, 2(6), 981-1022.

Author's Note

*The structural arguments presented in this paper draw on three independently developed frameworks: the grammar-entry of belief formation (developed by the author beginning in 2011); the directional drift of transmitted belief (developed in subsequent work); and the closure and remainder framework presented in *Consciousness, Closure, and the Cosmos*. These frameworks are presented here under descriptive names suited to the paper's argument. Their original designations and full development appear in the companion papers in which they are developed in full.*