Asian Journal of Philosophy (2023) 2:53
https://doi.org/10.1007/s44204-023-00104-z

ORIGINAL ARTICLE

What hinge epistemology and Bayesian epistemology can learn from each other

Olav Benjamin Vassend[1]

Received: 31 January 2023 / Accepted: 29 August 2023 / Published online: 27 September 2023

Abstract
Hinge epistemology and Bayesianism are two prominent approaches in contemporary epistemology, but the relationship between these approaches has not been systematically studied. This paper formalizes the central commitments of hinge epistemology in a Bayesian framework and argues for the following two theses: (1) many of the types of claims that are treated as paradigmatic hinges in the hinge epistemology literature, such as the claim that there exists an external world of physical objects, are not capable of enabling rational inquiry, even though this is typically regarded as a central property of hinges; (2) the standard Bayesian story of how rational inquiry proceeds is incorrect or at best incomplete.

Keywords Hinge epistemology · Bayesianism · Conditionalization

1 Introduction

In his last set of notes, published after his death as *On Certainty* (hereafter "OC"), Wittgenstein argues that there are special "hinge propositions" that any agent needs to take for granted when evaluating the plausibility of other, non-hinge propositions (Wittgenstein, 1969). Because of their foundational role in rational inquiry, Wittgenstein contends that hinge propositions are exempt from doubt. Taking Wittgenstein's remarks as their starting point, several authors have developed various versions of "hinge epistemology" in recent years (Coliva, 2015; Jenkins, 2007; Moyal-Sharrock, 2004; Pritchard, 2015; Salvatore, 2015; Volpe, 2021; Wright, 2004), mostly in a traditional epistemological framework that treats belief and evidential influence in a qualitative manner.

At the same time, formal methods have become increasingly popular among epistemologists. Chief among the various formal approaches is Bayesian epistemology,

✉ Olav Benjamin Vassend
olav.vassend@inn.no

[1] Department of Philosophy, Law and International Studies, Inland Norway University of Applied Sciences, Lillehammer, Norway

which uses probability theory to model confirmation, justification, testing, and rational belief change.[1] Although there have been a few attempts at formulating arguments from the hinge epistemology literature in a probabilistic framework (e.g., Moretti and Wright (2023), no one has so far given an in-depth treatment of the relationship between Bayesian epistemology and hinge epistemology. This paper aims to fill this gap. A major goal will be to formulate the various commitments of hinge epistemology in a Bayesian framework, with an eye toward seeing what the two approaches to epistemology can learn from each other. The main lessons that will be drawn are as follows: (1) if we take the central commitments of hinge epistemology seriously, then the standard Bayesian story of how rational inquiry proceeds is incorrect or at best incomplete; (2) the kinds of propositions that are often treated as paradigmatic hinges in the hinge epistemology literature, such as the claim that there exists an external world of physical objects, do not and cannot act as hinges that enable rational inquiry because they are too vague to play a meaningful role in probabilistic deliberation.

The rest of the paper is organized as follows. Section 2 lays out the core commitments of hinge epistemology. Section 3 introduces the core commitments of Bayesian epistemology. Section 4 provides a Bayesian formalization of the core commitments of hinge epistemology. Section 5 argues, contrariwise to what is commonly thought, that many paradigmatic examples of hinge propositions do not enable rational inquiry. Section 6 takes up the question of what assumptions do, in fact, enable rational inquiry. Section 7 draws larger philosophical implications for both hinge epistemology and Bayesian epistemology.

2 Core commitments of hinge epistemology

Hinge epistemology is a class of views rather than a single position, so the main goal of this section will be to isolate the core commitments that most hinge epistemologists are likely to accept, no matter what specific version of hinge epistemology they prefer. As was mentioned in the introduction, the central idea in hinge epistemology is that there are some propositions—the hinges—that are exempt from standard demands for justification because of the role they play in rational inquiry and evaluation. Standard examples of hinges range from specific claims such as "I have a brain" and "I have two hands" to vaguer claims such as that "the world is very old" and "there exists an external world of physical objects." What do all these propositions have in common?

Wittgenstein writes, "My having two hands is, in normal circumstances, as certain as anything I could produce in evidence for it… …That is why I am not in a position to take the sight of my hand as evidence for it" (OC §250). Thus, hinge propositions are special in at least two respects. First, they are certain—or at least more certain than any other contingent proposition. I take this to be a core commitment of hinge epistemology. Second, it is not possible to justify or question a hinge

[1] For a comprehensive introduction, see Titelbaum (2022a) and Titelbaum (2022b).

proposition by providing evidence for or against it. Wittgenstein writes, "…the questions that we raise and our doubts depend upon the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn" (OC, §341). Because hinge propositions are immune to doubt, hinge epistemology is often presented as a response to skepticism. A radical skeptic is someone who doubts whether we can know anything at all; a hinge epistemologist insists that all doubting presupposes a set of indubitable hinges that make the activity of doubting possible.

Although hinges are not subject to doubt, Wittgenstein maintains that the status of being a hinge is not absolute: a proposition can be a hinge at one time and a regular (doubtable and testable) proposition at another. Consider an example of an archetypical hinge proposition, such as "I have hands," Clearly, we can imagine circumstances in which this is an ordinary (and even false) empirical claim rather than an indubitable proposition. Wittgenstein himself writes: "It might be imagined that some propositions of the form of empirical propositions, were hardened and functioned as channels for such empirical propositions as were not hardened but fluid; and that the relation altered with time, in that fluid propositions hardened, and hardened ones become fluid… …The same proposition may get treated at one time as something to test by experience, at another as a rule of testing" (OC, §96-98). Whether *all* hinges can lose their hinge status is a matter of contention in the hinge epistemology literature. For example, Garavaso (1998) argues that Wittgenstein thinks all hinge propositions may be given up, whereas Moyal-Sharrock (2004) argues that there are certain "universal" hinges that cannot be given up because they form the "bounds of sense"—in particular, they need to be taken for granted in any sort of empirical inquiry. However, presumably all hinge epistemologists would have to agree that *most* hinges can lose their hinge status, given the right circumstances.

The quote by Wittgenstein in the preceding paragraph also contains what is arguably the most distinctive and important commitment of hinge epistemology: hinges are "rules of testing" that are necessary, and make it possible, for our perceptual evidence to bear on our beliefs about the empirical world, that is, hinges are propositions we need to take for granted in empirical inquiry. For example, suppose I seem to see my friend walking down the street. Then for me to be able to conclude that my friend, in fact, is walking down the street, I need to assume (or accept) the hinge proposition that the external world of physical objects exists. This is clear because if I instead assume that I am a brain in a vat, for example, then my sense impression of seeing my friend will not provide any support for my belief that my friend is walking down the street. On the other hand, if I do assume that the external world of physical objects exists, then it seems that my sense impression does indeed provide support for my belief.[2]

In summary, the core commitments of hinge epistemology, expressed informally, are as follows:

[2] Here I speak of sense impressions supporting beliefs, but it is worth flagging that Bayesian epistemology, which will be described in the next section, instead assumes that it is propositions (or sentences) that support propositions (or sentences).

(C1) Hinges are certain or at least more certain than other contingent propositions.
(C2) Hinges cannot be doubted.
(C3) (Most) hinges are not absolute: (most) hinges can lose the status of being a hinge; they can be given up.
(C4) Hinges are necessary, and make it possible, for our perceptual evidence to bear on our beliefs about the empirical world.

One of the main tasks of this paper will be to translate these informal commitments into a probabilistic framework; this is done in Section 4. But first, the next section introduces the basics of the Bayesian framework. Readers who are familiar with this framework may skip to Section 4 (although take note of the Criterion of Relevance, which will be used in later sections of the paper).

3 Core commitments of Bayesian epistemology

Bayesian epistemology is a framework for doing epistemology rather than a single philosophical position, and just as in the case of hinge epistemology, it comes in several varieties (Good, 1983). Hence, the goal of this section will be to identify the central commitments of the framework. To (subjective) Bayesian epistemologists, the basic entities of interest are the degrees of belief or plausibility assessments of a rational agent. Bayesians further assume, or sometimes try to demonstrate, that these plausibility assessments can be quantified by a numerical measure, $p()$, such that, for any propositions A and B, $p(A)$ represents the agent's degree of belief that A is true and $p(A \mid B)$ represents the agent's degree of belief that A is true given the assumption that B is true.

According to Bayesians, rational agents' degrees of belief obey several constraints. First, suppose we have an algebra of propositions, i.e., a set of propositions closed under conjunctions, disjunctions, and negations.[3] Then Bayesians hold that an agent's degrees of belief about the various propositions in the algebra should satisfy the following constraints, if the agent is rational (Titelbaum, 2022a, 2022b):

Probabilism:

1. If T is a necessary truth, then $p(T) = 1$.
2. $p(A)$ lies between 0 and 1, inclusive, for all propositions A.
3. If A and B are mutually exclusive, $p(A\ or\ B) = p(A) + p(B)$.[4]

Second, on the condition that $p(B) > 0$, Bayesians hold that an agent's conditional and unconditional degrees of belief should obey the following joint constraint:

$$\text{Ratio formula}: p(A \mid B) = p(A \& B)/p(B)$$

[3] This means that if A and B are in the set, then $A \& B$, $A\ or\ B$, $\sim A$, and $\sim B$ are all in the set.

[4] This is finite additivity, which is often replaced with a different assumption called "countable additivity." The differences between these two assumptions will not matter in this paper.

Here, $p(A \mid B)$ is the probability of A "conditional on B" or "conditioned on B." The ratio formula is often regarded as a definition of conditional probability, but following Hájek (2003), I think it is more plausible to treat both unconditional and conditional probability as primitive concepts and the ratio formula as a substantive normative constraint on how these two kinds of probability should relate to each other. One of the many reasons why we should think of conditional probability as primitive rather than as defined in terms of the ratio formula is that it is often easy to associate a number with $p(A \mid B)$ even if $p(B)=0$. For example, suppose you have a coin whose bias you are certain is 0.6 in favor of heads. Then p(The bias of the coin is 0.7) = 0, and yet, it is clearly the case that p(The coin comes up heads ∣ The bias of the coin is 0.7) = 0.7.

The ratio formula and the probability axioms have the following important consequences, whose proofs I omit here for the sake of brevity:

$$\text{Bayes}'\text{s formula}: p(A \mid B) = p(B \mid A)p(A)/p(B)$$

$$\text{Law of total probability}: p(A) = p(A \mid B)p(B) + p(A \mid \sim B)p(\sim B)$$

The probability axioms and the ratio formula are "synchronic" norms that say how degrees of belief at a given time should relate to each other. In addition to these synchronic norms, most Bayesians also accept the following learning norm for how degrees of belief should be updated in light of new evidence:

Bayesian conditionalization:

Suppose that an agent has a conditional degree of belief of $p(A \mid E)$ and that the agent learns E, and nothing else. Then the agent's unconditional degree of belief after learning the evidence, $p_E(A)$, should equal $p(A \mid E)$.[5]

I say that "most" Bayesians accept Bayesian conditionalization, because there are some Bayesians (broadly construed) who do not think it is the correct learning norm. One problem with Bayesian conditionalization is that if one conditionalizes on some proposition E, then E will henceforth have a probability of one, since $p_E(E)=p(E \mid E)=1$, but this consequence conflicts with another popular principle called "regularity," according to which no contingent proposition should every be assigned an extreme probability—I will discuss these issues again later in the paper. Suffice it to say here that, for these reasons, as well as other ones, some Bayesians instead prefer the following, more complicated learning norm, which reduces to Bayesian conditionalization when one becomes certain of the evidence:

Jeffrey conditionalization (simplified version adapted from Lin (2022)):

If the direct experiential impact on one's credences causes the credence in E to rise to a real number e (which might be less than 1), then one's credences should be changed as follows:

> For the possibilities inside E, rescale their credences upward by a common factor so that they sum to e; for the possibilities outside E, rescale their credences

[5] There are several versions of this norm (e.g., Pettigrew (2020) and Rescorla (2021)), but their differences will not matter here.

downward by a common factor so that they sum to 1. Reset the credence in each other proposition H by adding up the new credences in the possibilities inside H.

Probabilism, the ratio formula, and Bayesian conditionalization (or Jeffrey conditionalization) form the normative core of Bayesianism. The last ingredient we need to be able to formalize the commitments of hinge epistemology is not a norm but rather the following standard Bayesian explication of what it means for one proposition to support, disconfirm, or be irrelevant to another proposition:

Criterion of Relevance:

Given propositions A and B, B is positively relevant to (or confirms, or supports) A if $p(A \mid B) > p(A)$; B is negatively relevant (or disconfirms) A if $p(A \mid B) < p(A)$; and B is irrelevant to A if $p(A \mid B) = p(A)$.

The next section turns to the task of translating the commitments of hinge epistemology that we identified in Section 2 (i.e., (C1)–(C4)) into the core Bayesian framework presented in this section.

4 Formalizing the hinge commitments probabilistically

Let us start with commitment (C3), which states that (most) hinges are not absolute: (most) hinges can lose the status of being a hinge; they can be given up. Now, to a Bayesian, the notion that any sort of proposition can be "given up" is somewhat artificial, since Bayesians conceive of learning as a continuous process, where the probability of each proposition is adjusted up or down as new evidence comes in. In this context, to say that a proposition can be "given up" plausibly means that its probability can decrease when new evidence is considered. On this way of looking at things, a proposition that *cannot* be "given up" is plausibly a proposition whose probability cannot change at all, no matter what evidence is received. Are there any such propositions? The answer is "yes": if a proposition has a probability of one (or zero), then it follows from Bayesian conditionalization and the ratio formula that it is impossible for its probability to change via conditionalization no matter what evidence is received.[6] For similar reasons, it is also impossible for its probability to change via Jeffrey conditionalization. These considerations support the following formalization of (C3): (most) hinges do not have a probability of one. The preceding statement is qualified because, as we saw earlier, there are some hinge epistemologists who hold that there are certain hinge propositions that cannot be given up.

However, there are independent reasons for thinking that even hinge propositions that cannot be given up—if they exist—cannot have a probability of one. Recall that, according to (C4), hinges are supposed to be propositions that make it possible for perceptual evidence to bear on empirical propositions. Hence, the probability one assigns to a hinge proposition (if indeed hinges have probabilities at all, which we

[6] If $p(A) = 1$, then Bayesian conditionalization, Bayes' formula, and the law of total probability imply that $p_E(A) = p(A \mid E) = p(E \mid A)p(A)/(p(E \mid A)p(A) + p(E \mid {\sim}A)p({\sim}A)) = p(E \mid A)/p(E \mid A) = 1$.

will soon see there are good reasons to doubt) is supposed to be assigned independently of—and prior to—one's consideration of the empirical evidence. Or, at least, if H is a hinge proposition that we need to take for granted for it to be possible for us to evaluate the plausibility of empirical proposition Q, then whatever probability we assign to H needs to be assigned independently of Q.

But consider a purported universal hinge such as "there exists an external world of physical objects." One reason why some hinge epistemologists hold that this is a hinge that cannot be given up is because it (supposedly) needs to be taken for granted in any sort of empirical investigation. But if "there exists an external world of physical objects" must be assumed in order for any sort of empirical investigation to be possible, then any probability one gives to this proposition must presumably be assigned independently of all possible empirical evidence—i.e., the probability must be assigned purely on a priori grounds.

In the Bayesian literature, the standard examples of propositions that have a probability of one, independently of any empirical evidence, are propositions that are necessarily true, where the necessity in question is logical or perhaps mathematical. However, archetypical hinges like "there exists an external world of physical objects," "the world has existed for a very long time," and "no one has ever been to the moon" are straightforwardly empirical claims that are merely contingently true, and according to the aforementioned regularity principle, contingent propositions should never be assigned extreme probabilities.

Regularity is not a core commitment of Bayesianism, but it has a compelling justification: as we saw earlier, in a Bayesian framework, it is possible to learn that P is false only if P is assigned a non-extreme probability. A proposition that is contingent is a proposition that *could* be false, and it seems irrational to deliberately put ourselves in a position where we make it impossible for ourselves to learn that some proposition is false, even though we recognize that it might be false. Hence, no contingent proposition—including hinges—should be assigned an extreme probability.

An objection to the regularity principle is that if we use Bayesian conditionalization, then any proposition E that has already been conditioned on must have a probability of one, since $p_E(E) = p(E \mid E) = 1$. Hence, the objection goes, an agent who updates on evidence via conditionalization is bound to violate regularity, which implies that regularity cannot be correct.[7]

One rather revisionary response to this objection is to keep the regularity principle and instead give up on Bayesian conditionalization, perhaps by replacing it with Jeffrey conditionalization. I think this response has much going for it, especially since there are independent grounds for preferring Jeffrey conditionalization.[8]

But even if we decide to stick with standard Bayesian conditionalization, that does not mean we should reject regularity wholesale. Although it is true that Bayesian conditionalization forces us to assign a probability of one to propositions that we have *already learned*, the motivation behind following the regularity principle is

[7] I am grateful to a referee for suggesting this objection.

[8] For example, that it gives an account of how learning proceeds when an agent learns a piece of evidence without thereby becoming certain that the evidence is true.

arguably that we do not want to decisively bias inquiry concerning propositions we have *yet to learn*. We can therefore formulate the following weaker and more plausible version of regularity, which is consistent with Bayesian conditionalization:

Moderate regularity: if *P* is a proposition that is contingent and that we have not yet learned (by conditionalizing on it), then we should not assign *P* an extreme probability.

Whatever epistemic status hinge propositions have, it is probably fair to say that hinge epistemologists do not think of them as propositions that we can learn through a process such as conditionalization. Indeed, hinge propositions are supposed to be assumptions that enable us to undertake empirical investigation in the first place, and so cannot—on pain of circularity—themselves be learned through empirical investigation. If it is correct that hinges cannot be learned through conditionalization, then it follows from the moderate version of regularity that hinges should not have a probability of one.

An entirely independent reason for thinking that hinge propositions cannot or should not have a probability of one has to do with the connection that many Bayesians draw between probability and betting odds.[9] According to many Bayesians, the probability one assigns to a proposition is supposed to reflect the amount of money one would be willing to wager that the proposition is true. On this view, assigning a probability of one to a proposition would entail that one would and should be willing to wager any arbitrarily large amount of money on a bet that pays out $1 if the proposition is true. I am not sure what amount of money I would be willing to wager on the proposition that there exists an external world of physical objects (or that no one has been to Mars), but at least in my case, it is not any arbitrarily large amount.

For all the above reasons, then, we arrive at the following formalization of (C3), which in fact is stronger than the informal version of the claim, since it is not qualified: hinges are not propositions that have a probability of one.

Our formalization of (C3) has immediate implications for how we ought to formalize (C1) as well. In a Bayesian context, it is standard to identify certainty with maximal probability, so that a proposition is certain if and only if it has a probability of one. However, in light of the preceding discussion, this is not a plausible way of understanding what it means for a hinge proposition to be certain or more certain than other propositions.

Before we make a different attempt at formalizing (C1), let us have a closer look at (C4). (C4) says that hinges are necessary, and make it possible, for our perceptual evidence to bear on our beliefs about the empirical world. There is both a necessity claim and a possibility claim here, and it makes sense to treat these claims separately. Let us start with the necessity claim—how should we best understand this claim from a Bayesian perspective? As I understand it, part of the claim is that hinges are required for it to be possible to make a rational plausibility assessment of any (non-hinge) empirical claim. For example, in order for me to be in a position to rationally assign a probability to the proposition that it is going to rain tomorrow, I need to make the hinge assumption that (say)

9 I am grateful to a referee's comment for inspiring this paragraph.

there is an external world of physical objects, such as rain, in the first place. Or, to put the point more formally (and indeed a bit more weakly), part of the necessity claim is that in order to come up with an assessment of the plausibility of an empirical (non-hinge) proposition *B*, I must first consider how plausible *B* is in light of some suitable hinge proposition *H*—i.e., the conditional probability $p(B \mid H)$ is epistemically prior to the unconditional probability $p(B)$ (assuming empirical propositions can have unconditional probabilities at all, which we will soon see there are reasons to doubt).

However, the necessity claim arguably contains a second part as well, which we can formalize in terms of the Criterion of Relevance from Section 3. Recall that this criterion says that *B* is relevant to *A* if and only if $p(A \mid B) \neq p(A)$. Now, we saw earlier that in order for evidence to have an influence on the probability of some proposition, it is a necessary condition that the proposition not have a probability of zero (or one). Hence, the claim that some hinge *H* is necessary in order for evidence to bear on some proposition *B* can be naturally understood as saying that the probability of *B* is non-zero only if we assume *H*; in other words, conditioning on ~*H* gives *B* a probability of zero. Summing up the above discussion, (C4) plausibly states that hinges have the following property:

Necessity Property:

If *H* is a hinge for some proposition *B*, then the following holds:

(1) The conditional probability $p(B \mid H)$ is epistemically prior to $p(B)$ in the sense that a rational agent cannot have the latter degree of belief without having the former.

$$p(B \mid \sim H) = 0.$$

In what follows, I will say that *H* has the Necessity Property with respect to *B* if and only if the above condition holds. For example, "there is an external world of physical objects" has the Necessity Property with respect to my belief that I am seeing my friend walking toward me, and this is reflected in the fact that I (plausibly) cannot evaluate how plausible it is that my friend is walking toward me, all things considered, without first evaluating how plausible it is given the assumption that there exists an external world of physical objects; and, on the other hand, if I do make the assumption that there is no external world of physical objects, then:

p(My friend is walking towards me | There is no external world) = 0.

In addition to stating that hinges are necessary, (C4) also says that hinges make it possible for our evidence to bear on our beliefs; hinges enable rational evaluation. To a Bayesian, a piece of evidence is (plausibly) a proposition that it is possible to learn via conditionalization. Thus, the natural way of understanding this claim in terms of the Criterion of Relevance is as saying that if *H* is a hinge for some proposition *B*, then there exists some piece of (perceptual) evidence *E*, such that *E* can be learned through conditionalizing and such that conditioning the probability distribution on *H* renders *E* relevant to *B*. More formally:

Enabling Property:

If H is a hinge proposition for some proposition B, then there exists evidence E, such that it is (in principle) possible to learn E through conditionalization and such that $p(B \mid H) \neq p(B \mid E, H)$.

In what follows, I will say that "H has the Enabling Property with respect to B" if and only if there exists evidence E such that conditioning on H renders E relevant to B. For example, "there is an external world of physical objects" seems to have the Enabling Property with respect to my belief that my friend is walking toward me because it seems to be the case that:

p(My friend is walking towards me | I seem to see my friend walking towards me & There is an external world of physical objects) >p(My friend is walking towards me | There is an external world of physical objects).[10]

We are finally able to capture what (C4) says in a Bayesian framework as follows: hinges are propositions that have both the Enabling Property and the Necessity Property with respect to at least one proposition. Note that this does not mean that having the Enabling Property and the Necessity Property is a necessary *and sufficient* condition for a proposition to be a hinge, as there are examples of propositions that have both properties and yet do not seem to be hinges.[11] However, if H is a hinge proposition and H has both the Necessity and Enabling Property with respect to some proposition B, I will say that H is a hinge for B.

The Necessity Property can also help us make sense of the claim that hinges are more certain than other propositions. In particular, the Necessity Property immediately implies that hinges have the following property:

Maximal Plausibility: for any hinge proposition, if the hinge proposition were to be assigned a probability, then its probability would be at least as high as the probability of any proposition for which it is a hinge.

Maximal Plausibility follows from the Necessity Property because if H is a hinge for some proposition Q and H is assigned a probability, then: $p(Q) = p(Q \mid H)p(H) + p(Q \mid {\sim}H)p({\sim}H) = p(Q \mid H)p(H) \leq p(H)$.

Because any existential claim about the empirical world seems to entail or at least strongly support very general claims of the form "there is an external world" and "the world did not just pop into existence," it seems that general hinges have the following property:

General hinge property: For any non-hinge existential claim Q about the empirical world and any (sufficiently) general hinge H, $p(Q \mid {\sim}H) = 0$.

It follows from Maximal Plausibility and the general hinge property that general hinges, such as "there exists an external world," have a higher probability than all non-hinge existential empirical claims, given that hinge propositions have probabilities in the first place. Maximal Plausibility therefore captures the sense in which hinges are more certain than other propositions and is consequently a natural way of understanding (C1) in a Bayesian framework.

[10] In the next section, I will dispute that this inequality holds, but prima facie it seems plausible.

[11] I am grateful to two referees for both pointing this out and for providing plausible examples of propositions that have both the Necessity and Enabling property and yet are plausibly not hinges. For the sake of brevity, I omit providing any specific examples here.

Let us now turn to (C2), which states that hinges cannot be doubted. How are we to understand this claim? Clearly, anything "can" be doubted in the sense that, for any proposition Q that you might decide to assert, I can sincerely utter the phrase "I doubt Q." So, the question is not whether I can doubt Q in the weak sense that I can form an attitude of doubt toward it, but whether I can have legitimate (or rational) grounds for doubting Q. Charitably construed, (C2) says that there cannot be legitimate (or rational) grounds for doubting a hinge proposition.

In a Bayesian framework, doubting a proposition Q plausibly involves finding some proposition D such that D, if true, would lower the probability of Q, that is, D is grounds for doubting Q only if $p(Q \mid D) < p(Q)$. This probabilistic inequality is a necessary but not sufficient condition; additional constraints must generally be placed on D for it count as a legitimate ground for doubt. For example, the mere possibility that $\sim Q$ might be true is not a legitimate ground for doubting Q. I will not attempt to set out a set of necessary and sufficient conditions for what makes some proposition D a legitimate ground for doubting some other proposition Q. However, it is plausible that if Q is an empirical proposition (as hinge propositions typically are), then a legitimate ground for doubting Q must itself be an empirical proposition. This does not necessarily mean that the grounds for doubt, D, must be learnable through conditionalization, but it does mean that D must be a proposition whose plausibility could in principle change in light of evidence that we could learn through conditionalization (or Jeffrey conditionalization). For example, "it will rain tomorrow" is not a proposition we can learn through conditionalization (at least not now), but it is still empirically investigable because we can assess its plausibility now in light of evidence that we can learn through conditionalization.

Whether a proposition is empirically investigable in this sense may be historically contingent, of course, in the same way that it may (at least according to some hinge epistemologists) be historically contingent whether a proposition has hinge status. However, as we discussed earlier, due to the constitutive role that hinge propositions are supposed to play in empirical inquiry, they are arguably not themselves the sort of proposition that it is possible to investigate through empirical investigation, at least so long as they are hinges. At the very least, then, it follows that the set of propositions that are legitimate grounds for doubting a hinge proposition H does not include $\sim H$ or propositions entailed by $\sim H$, because if $\sim H$ can be investigated through empirical investigation, then so can H (the probability axioms guarantee that an assessment of the plausibility of $\sim H$ is automatically also an assessment of H).

Summing up the discussion so far, we arrive at the following formalization of (C2): for every hinge H, there does not exist a proposition D, such that D in principle could be investigated empirically and such that $p(H \mid D) < p(H)$. Note that this immediately implies that there also does not exist a proposition S, meeting the same conditions, but such that $p(H \mid S) > p(H)$. This is because if $p(H \mid S) > p(H)$, then $p(H \mid \sim S) < p(H)$.[12] In other words, on this proposal, for every

[12] If $p(H \mid D) < p(H)$, then $p(D \mid H)p(H)/p(D) < p(H)$ and hence $p(D \mid H) < p(D)$, which implies that $p(\sim D \mid H) > p(\sim D)$, and hence $p(\sim D \mid H)p(H)/p(\sim D) > p(H)$, which finally implies that $p(H \mid \sim D) > p(H)$.

hinge H and for all propositions Q such that Q in principle could be empirically investigated, $p(H \mid Q) = p(H)$, i.e., hinges are evidentially irrelevant to all propositions that could in principle be empirically investigated.

However, in light of our earlier discussion of (C3) and (C4), the claim that hinges are evidentially irrelevant to all empirically investigable propositions cannot be correct. Suppose H is a hinge for some empirical claim Q and that H has a probability less than one. Then the Necessity Property implies that $p(Q \mid \sim H) = 0$, which in turn implies that $p(Q) = p(Q \mid H)p(H)$ (because we are assuming that H has a probability less than one). But then, $p(H \mid Q) = p(Q \mid H)p(H)/p(Q) = 1 \neq p(H)$. So, if H is a hinge for Q and H does not have a probability of one, Q must be evidentially relevant to H.

In other words, any proposition H that is a hinge for some empirically investigable proposition Q also opens itself to being doubted on the basis of Q, so long as H has a non-extreme probability. For example, if "there is an external world" is a hinge for "my friend is walking toward me" and the plausibility of "my friend is walking toward me" changes, then, according to the above discussion, the probability that there is an external world will have to change as well. Hence, if we assume that hinges do not have (unconditional) probabilities of one, which we have seen there are strong reasons to accept, then it seems it cannot both be true that hinges have the Necessity Property and that hinges are not subject to doubt. To sum up, what we have shown is that the following claims are jointly inconsistent:

(a) Hinges have unconditional probabilities.
(b) For every hinge H, it is not the case that $p(H) = 1$.
(c) Hinges have the Necessity Property.
(d) For every hinge H, it is not the case that there exists a proposition, D, such that D could in principle be learned through empirical investigation and such that $p(H \mid D) < p(H)$.

We have already seen that there are strong reasons to accept (b) and (c), and (d) is a straightforward formalization of the claim that hinges cannot be doubted. It follows that we must reject (a) and conclude that hinges do not have unconditional probabilities. Of course, in the Bayesian framework, any proposition that does not have an unconditional probability cannot be doubted, cannot be tested, and cannot be justified, since doubting, testing, and justification are all conceived of probabilistically (ultimately in terms of the Criterion of Relevance). So, the idea that hinges do not have unconditional probabilities fits well with the idea that hinges cannot be doubted.

Incidentally, Moretti and Wright (2023)also propose that hinges do not have (unconditional) probabilities as a solution to the perplexing finding that, given a few assumptions, the confidence one has in a hinge based on perceptual evidence cannot be greater than the confidence one has that one is not in a skeptical scenario. Moretti and Wright argue that there are only two viable responses to this problem: either hinges have a probability of one or they do not have probabilities at all. Our discussion provides additional, and independent, support to the view that hinges do not have unconditional probabilities.

The conclusion that hinges do not have unconditional probabilities does not imply that they cannot have *conditional* probabilities, of course. Indeed, even if we cannot (rationally) assign an unconditional probability to H, we can clearly still say that $p(H \mid H) = 1$, for example. Furthermore, we saw earlier that if H is a hinge for Q and H has an unconditional probability, then $p(H \mid Q) = 1$. The mathematical argument for this conclusion depended on the assumption that H has an unconditional probability, but, arguably, it is independently plausible that $p(H \mid Q) = 1$, for if H is a hinge for Q, then one considers H a necessary assumption in order for Q to be true, and so on the condition that Q is true, one should be certain that H is true as well. But if $p(H \mid Q) = 1$, then Bayes's formula says that $1 = p(H \mid Q) = p(Q \mid H)p(H)/p(Q)$, which implies that $p(H) = p(Q)/p(Q \mid H)$. Thus, if Q has an unconditional probability, it seems H will automatically be forced to have an unconditional probability as well. The upshot, then, is that Q cannot have an unconditional probability either given that H does not.

A further upshot is that if H is a hinge for Q, then it cannot be possible to learn Q via conditionalization. This is because if $p(H \mid Q) = 1$, then conditionalizing on Q would entail that H received an unconditional probability of 1, which is contrary to our conclusion that H (as long as it is a hinge) cannot have an unconditional probability.[13]

The preceding conclusions may seem radical, but they are arguably in line with the spirit behind hinge epistemology. According to hinge epistemologists, hinges are propositions that are necessary for us to be able to bring our perceptual evidence to bear on our empirical beliefs. For example, without the hinge assumption that there is an external world, the proposition that I have a visual impression of seeing my friend down the street cannot provide probabilistic support to the proposition that my friend in fact is down the street. It stands to reason, then, that once I conditionalize on my visual impression, my degree of belief that my friend is down the street will still be conditional on the hinge assumption that there is an external world of physical objects.

Summing up, we have arrived at the following way of formalizing the commitments of hinge epistemology in the Bayesian framework:

(C1) Hinges have the property of Maximal Plausibility.
(C2) Hinges do not have probabilities.
(C3) Hinges are not propositions that have a probability of one.[14]
(C4) Hinges are propositions that have both the Enabling Property and the Necessity Property with respect to at least one proposition.

The goal of the remainder of the paper is to use the formalization provided in this section to draw important conclusions for both hinge epistemology and Bayesian epistemology.

[13] I am grateful to a reviewer for helping me see all these consequences.

[14] Of course, given our formalization of (C2), (C3) becomes redundant. But note that (C3) was crucial in the argument that established (C2).

5 Why general propositions do not and cannot enable rational inquiry

Much of the attention in the hinge epistemology literature has focused on purported universal hinges, such as “there exists an external world,” “the Earth is very old,” and even “I am not massively mistaken about everything.” The reason is probably that much of the discussion has been concerned with rebutting skepticism, and these sorts of claims are exactly what the radical skeptic denies we can know. The linchpin premise in the hinge epistemologist argument against skepticism is the claim that hinges have property (C4): they are necessary for, and make possible, rational evaluation of other propositions. This property is what gives hinges their special status and supposedly makes them impervious to skeptical doubt.[15]

For example, according to many hinge epistemologists, we need to assume (or accept, or otherwise take for granted) something along the lines of the claim that “the Earth is very old” for it to be possible for geological evidence to bear on more specific claims about the age of the Earth, such as the claim that the Earth is 4.5 billion years old, and this is what makes it rational (or entitles us) to accept the claim that the world is very old, even before taking into consideration specific evidence concerning the age of the Earth. However, the aim of this section is to argue that when viewed through a Bayesian lens, it becomes clear that these kinds of very vague claims do not and cannot play a meaningful role in rational inquiry.

Let us consider more closely the case of the purported hinge, “the Earth is very old.” If this proposition is a hinge for the claim that the Earth is 4.5 billion years old, then the Enabling Property entails that there must exist evidence E such that:

$$p(\text{Earth is 4.5 billion years old} \mid \text{Earth is very old}) \neq p(\text{Earth is 4.5 billion years old} \mid \text{E\&Earth is very old})$$

Let us consider the first of these conditional probabilities. What is the probability that the Earth is 4.5 billion years old given that the Earth is very old? In trying to answer this question, we immediately run into problems because it is quite clear that whether something is “very old” is contextual and depends on whether we are talking about insects, elephants, societies, or planets. This point is obvious, but it has important implications. Presumably, people 500 years ago would have agreed that the Earth is “very old” if we had asked them. However, if we pressed them to be more particular about what they have in mind, it is plausible that they (or at least many of them) would have clarified that they think the Earth is on the order of a few thousand years old. But note that for any piece of evidence E, we have:

$$p(\text{Earth is 4.5 billion years old} \mid \text{Earth is a few thousand years old})$$
$$= p(\text{Earth is 4.5 billion years old} \mid \text{E\&Earth is a few thousand years old}) = 0$$

[15] In the preceding section, (C4) also turned out to be the most central hinge commitment, since the informal version of (C4) was important in establishing the claim that no hinge can have a probability of one, and this claim, in turn, was crucial for the argument that hinges do not have probabilities at all and therefore cannot be doubted.

In other words, if we assume that “the Earth is very old” means that the Earth is on the order of a few thousand years old, which is plausibly what many people 500 years ago thought, then “the Earth is very old” cannot have the Enabling Property with respect to the claim that the Earth is 4.5 billion years old. This, in turn, implies that people 500 years ago would have been incapable of appreciating the evidence for the geological age of the Earth if (hypothetically) they had been presented with it, simply because they had the wrong hinge assumption. Perhaps this is a conclusion that some people would be willing to accept, but then, the next question is why we today have a hinge assumption (namely, that the Earth is on the order of a few billion years old) that enables us—as opposed to people 500 years ago—to appreciate the evidence for the more specific claim that the Earth is 4.5 billion years old.

To get around these issues, one might maintain that when people 500 years ago held that the “The Earth is very old,” what they had in mind was instead that it is *at least* a thousand years old (or something along these lines), whereas what we have in mind today is that the Earth is *at least* a billion years. Clearly, it is possible for the following inequality to hold, and so, on this reading, it appears that “the Earth is very old” can have the Enabling Property after all, even for people 500 years ago:[16]

$$p\,(\text{Earth is 4.5 billion years old} \mid \text{Earth is a few thousand years old})$$
$$< p\,(\text{Earth is 4.5 billion years old} \mid \text{E\&Earth is a few thousand years old})$$

However, this suggestion faces at least two problems. First, much like the claim that the Earth is very old, the proposition that the Earth is at least a thousand years old is rather vague. Whether it is plausible that the Earth is 4.5 billion years old given the assumption that the Earth is at least a thousand years old depends crucially on further assumptions about *how much* older than a thousand years it is plausible that the Earth is. By itself, the assumption that the Earth is at least a thousand years old is too unspecific to enable a probability assignment to the proposition that the Earth is 4.5 billion years.

Second, and more importantly, this proposal does not address the more fundamental question of why people 500 years ago accepted or assumed that the Earth is at least a thousand years old, whereas we today accept or assume the more specific proposition that the Earth is at least a billion years. How did we all move from having the first assumption to having the second one?

From a Bayesian point of view, the natural way to model the situation is by supposing that each person (or at least each person who cares about the matter) has a prior probability distribution over all possible ages of the Earth, stretching from zero to infinity. Five hundred years ago, reasonable people presumably had probability distributions that were concentrated in the region from roughly 1000 to 10,000 years. But, at the same time, most people did not exclude the possibility that the Earth could be much older. Thus, they assigned some non-zero probability to the possibility that the age of the Earth was billions or even trillions of years (or infinite). As the evidence started coming in that the Earth in fact is significantly older than people had anticipated, their probability distributions shifted upward. Today,

[16] I am grateful to a referee for suggesting this possibility.

our probability distribution is concentrated in the region of 1 to 10 billion years (let us say), but we also assign some probability to the possibility that the Earth is much older or much younger.

In the above Bayesian story, there is no room for a pre-empirical vague hinge such as "the Earth is very old"; this assumption, if it is an assumption, simply plays no role in setting the relevant conditional probabilities. Rather, it is plausible that our understanding of the claim that the Earth is very old has changed over time in response to the empirical evidence. Five hundred years ago, when people thought (or assumed) that the Earth is very old, what most of them plausibly had in mind was that it is a few thousand years old; today, what most of us have in mind is that the Earth is a few billion years old.

Much the same can be said for the claim that there is an external world of physical objects. A claim of this sort is too vague to be able to figure in probabilistic reasoning and so rather than act as a pre-empirical hinge that allows us to connect perceptual evidence to specific empirical beliefs, it seems more likely that the flow of information has been in the opposite direction: that is, in the same way that our understanding of the claim that the Earth is very old has changed as we have learned more about the age of the Earth, it is plausible that our understanding of the claim that there exists an external world of physical objects has changed as our scientific knowledge of the world has grown and our understanding of the physical world has shifted. Today, when we say that there exists an external world of physical objects, what we have in mind is that there are atoms, quarks, electrons, forces, etc., or at least that the story told by modern science is roughly correct, although we may be open to the details being wrong. This is plausibly very different from what people 500 years ago had in mind when they maintained (or assumed) that there is an external world of physical objects.

Some readers may object to the idea that we have in mind something different today than what people did 500 years ago when we accept or assume that there exists an external world of physical objects. I am not going to defend this idea in detail, although I think it is plausible, but suppose, for the sake of argument, that we grant that we have in mind precisely the same proposition today that people had 500 years ago when we maintain that there exists an external world of physical objects. Given how much our understanding of the physical world has changed, it is clear that this assumption—whatever it is—must be very abstract and vague. And to the extent that this proposition is very vague and abstract, it is implausible that it can play a meaningful role in probabilistic deliberation. For example, how plausible it is that an apparition in front of me is an evil spirit given the assumption that "there exists an external world of physical objects" depends crucially on whether the class of physical objects contains evil spirits; similarly, how plausible it is that our instrument is detecting an electron given the assumption that there are physical objects depends on whether electrons are in the class of physical objects. Thus, again, the bare assumption that there is an external world of physical objects in some generic and vague sense cannot plausibly play an important role in probabilistic deliberation.

Other readers may agree with the idea that the assumption that there is an external world of physical objects is different today from what it was 500 years ago and yet maintain that this assumption—in the form it has today—can serve as an important hinge assumption that enables empirical inquiry *today*.[17] Thus, if we take "there exists an external world of physical objects" today to mean the same thing as "the world is roughly as described by modern science," then it seems that this proposition can perhaps play some role in probabilistic deliberation (at least if it is made somewhat more precise what is included in modern science). However, the idea that "the world is roughly as described by modern science" is a pre-empirical hinge that is necessary for, and enables, empirical (scientific) inquiry is implausible. Again, it is far more plausible that this is a proposition we accept today because of all the incremental (and large) successes that the scientific enterprise has had over the past several hundred years.

The conclusion, then, is that vague claims such as "the Earth is very old" and "there exists an external world of physical objects" cannot function as enabling hinges in rational inquiry, at least if we conceive of rational inquiry in a Bayesian way. They cannot be enabling hinges because they are too vague to play a meaningful role in setting the conditional probabilities that are required in Bayesian inference, and indeed it is plausible that the way we understand these claims has changed in direct response to the growth of more specific, scientific knowledge.

6 Which assumptions enable rational inquiry?

The fact that vague claims do not function as hinges for rational inquiry does not mean that there are no hinges. When scientists inferred that the Earth is around 4.5 billion years old, they clearly had to make assumptions. Indeed, it is a familiar point from philosophy of science that scientific inference almost always requires that one make "auxiliary assumptions" that connect evidence and hypotheses of interest,[18] and it is natural to think of such auxiliary assumptions as hinges of a sort. Usually, the auxiliary assumptions that occur in scientific inferences will be highly specific (often quantifiable) assumptions about the processes that produced the available evidence.

It might be useful to look at a simple example. Consider an attempt at estimating the mass of an object by repeatedly weighing the object on a scale. The results of the weighing do not, by themselves, say anything about the mass of the object unless we make assumptions about how the measurement outcomes relate to the object's mass. A standard way of connecting hypotheses of interest to measurement outcomes is by way of a statistical model. For example, we may assume that the measurements are

[17] I am grateful to a referee for pushing me to think much harder about these issues.

[18] This is the "Quine-Duhem" thesis (Duhem, 1998; Quine, 1951).Quine, W. V. (1951). Two Dogmas of Empiricism. *Philosophical Review*, *60*(1), 20-43. , ibid.

distributed identically and independently according to a normal distribution that is centered on the object's mass. With these assumptions in place, it becomes possible to assess the conditional probability of the measurement outcome given various possible values of the object's mass, $p(\text{outcome} = o \mid \text{mass} = m)$, and this in turn makes it possible to estimate the value of the object's mass.

The above analysis rests on several auxiliary assumptions or enabling "hinges," i.e., propositions that (at least during the estimation of the object's mass) are accepted and not tested, such as the assumption that the measurements are independent and distributed normally. The fact that these assumptions are not tested during the estimation process does not mean that they are not tested at all, however. For example, the assumption of normality may be assessed with a chi-squared test. Indeed, most standard statistical assumptions can be tested, at least in principle. Crucially, however, any additional test will in turn make further assumptions that will need to be taken for granted during the testing procedure. All this is in line with Wittgenstein's remark that "The same proposition may get treated at one time as something to test by experience, at another as a rule of testing" (OC, 96-98), if we interpret "rule of testing" to mean an assumption that is held fixed during the test. However, at some point, sooner or later—in practice, usually sooner—the testing will need to stop, and hence, there are inevitably assumptions that remain untested, assuming we do not accept circular testing, i.e., assuming we do not think it is legitimate to, for example, use assumption A_1 to test assumption A_2 and then use A_2 to test A_1. Hence, as a logical consequence of the fact that all testing requires assumptions and we do not accept as valid circular testing schemes, there will always be assumptions that we must take for granted without testing, assumptions that make it possible for our evidence to bear on the hypotheses of interest. These are the sorts of hinges on which statistical practice turn. Indeed, even in a thorough-going Bayesian statistical analysis, it is inevitably the case that the analysis rests on important assumptions that are not themselves assigned probabilities.[19]

But the sorts of assumptions that occur in statistical practice plausibly fail to meet another feature that hinges are supposed to have, namely that they are *necessary* for the rational evaluation of propositions to be possible. According to the Necessity Property, if H is a hinge for B, then $p(B \mid {\sim}H) = 0$. Indeed, as we saw earlier, what makes claims such as "There exists an external world of physical objects" special is that they seem to have the Necessity Property with respect to many of our more specific empirical beliefs. On the other hand, the sorts of assumptions that occur in statistical inference rarely or never have the Necessity Property with respect to the hypotheses of interest. For example, for measurement outcomes to bear on the various hypotheses about the mass of some object, it is not *necessary* to assume that the measurements are normally distributed; indeed, an infinite number of alternative assumptions would equally have enabled us to estimate the object's mass given the observations. Some assumption must be made, but no particular assumption is necessary or indeed privileged a priori over any other assumption.

[19] For example, Bayesian statistical analyses typically rest on exchangeability assumptions, and I have never seen any statistician assign a probability to an exchangeability assumption in a statistical analysis.

The fact that assumptions need to be made in the context of any empirical testing (as a logical necessity) and that no set of assumptions is universal or privileged a priori does not mean that there is a free-for-all when it comes to making suitable hinge assumptions, however. Clearly, some assumptions are (in some sense) more reasonable than others. How hinge assumptions are justified—how they gain their hinge status—is therefore an important question for hinge epistemology that is beyond the scope of this paper. As we will see in the next section, it is also beyond the scope of the typical Bayesian story of how learning, justification, and testing take place.

7 Philosophical implications for Bayesian epistemology and hinge epistemology

In Section 4, we concluded that hinges do not have unconditional probabilities. This conclusion is important for both hinge epistemology and Bayesian epistemology. For hinge epistemologists, it implies that hinges are propositions toward which we do not have (rational) degrees of belief, which in turn vindicates the view that our attitude toward hinges must be something other than a belief. If this is true, and if knowledge entails belief and belief is construed as coming in degrees, then it also follows that we cannot know hinge propositions.

The conclusion that hinges do not have probabilities has perhaps even greater consequences for Bayesianism. According to what is arguably the standard view among Bayesians, rational inquiry proceeds as follows: first, any proposition that the agent cares about should receive (or have) an unconditional probability and a probability conditional on the various pieces of evidence the agent thinks she may receive. Note that this does not mean that the agent should associate a probability with every conceivable proposition. However, at the very least, every proposition that plays a crucial role in the agent's deliberation should have an associated probability. Second, once evidence is received, the agent should conditionalize the probability of each of these propositions on the evidence.

Hinge propositions obviously play a crucial role in inference since they enable rational inquiry, so agents clearly care about hinge propositions. Yet, as we have seen, because of the role they play in rational inquiry, hinge propositions necessarily do not have unconditional probabilities. Hence, the standard Bayesian story is incorrect or at best incomplete since it leaves out an account of the role in rational inquiry played by hinge propositions.

Furthermore, one might have hoped that there is a single set of universal assumptions that undergird all rational inquiry, but the discussion in Section 6 indicates that this is not the case. No single set of hinge assumptions is necessary, and hence the set of enabling hinge assumptions is, in a sense, arbitrary and yet at the same time needs to be chosen judiciously, since some hinge assumptions clearly are more reasonable than others. But since hinge assumptions do not have probabilities, the question of how they ought to be justified and chosen must presumably come from outside of the core Bayesian framework, since the Bayesian framework treats all justification in terms of probability raising (in accordance with the Criterion of Relevance).

Another, very different, implication for Bayesians has to do with the aforementioned connection that some Bayesians draw between degrees of belief and betting dispositions. If hinge propositions do not have probabilities, then it follows that agents also cannot (or do not) have betting dispositions toward hinge propositions. Some Bayesians may be comfortable with accepting this consequence; others might revise their assumption that degrees of belief are necessarily connected to betting dispositions; of course, others yet might view this consequence as a reason to reject hinge epistemology.

The discussion in the preceding two sections also has an important implication for hinge epistemology: according to the argument in Section 5, many paradigmatic examples of hinges, such as "there exists an external world," cannot actually enable rational inquiry because they are too vague to enter meaningfully into probabilistic deliberation. If having the Enabling Property is a necessary condition for being a hinge, which I argued in Section 4 is the case, it follows that these sorts of claims are not actually hinges after all. On the other hand, in Section 6, we saw that the assumptions that arguably do enable rational inquiry (more particularly, statistical inquiry) do not have the Necessity Property. Hence, if the Necessity Property is a necessary property of hinges, then these assumptions are not hinges either. In fact, it is hard to come up with an example of any proposition that has both the Necessity Property and the Enabling Property. As I see it, there are two viable responses: either there are no hinges at all or hinges do not have the Necessity Property with respect to the propositions for which they act as a hinge. Both these responses are likely to have detrimental effects on the standard responses that hinge epistemologists give to the skeptic. However, I leave further discussion of this issue to a different occasion.

Funding Open access funding provided by Inland Norway University Of Applied Sciences

Data availability This paper has no associated data.

Declarations

Competing interests The author declares that he has no competing interests.

References

Coliva, A. (2015). *Extended rationality: A hinge epistemology*. London, UK: Palgrave Macmillan.

Duhem, P. (1998). Physical theory and experiment. In Martin Curd & J. A. Cover (Eds.), *Philosophy of Science: The Central Issues* (pp. 257–279). New York, NY: Norton.

Garavaso, P. (1998). The distinction between the logical and the empirical in on certainty. *Philosophical Investigations, 21*(3), 251–267. https://doi.org/10.1111/1467-9205.00070

Good, I. J. (1983). 46656 varieties of Bayesians. In M. N. Minneapolis (Ed.), *Good thinking: The foundations of probability and its applications* (pp. 20–21). University of Minnesota Press.

Hájek, A. (2003). What conditional probability could not be. *Synthese, 137*(3), 273–323.

Jenkins, C. S. (2007). Entitlement and rationality. *Synthese, 157*(1), 25–45. https://doi.org/10.1007/s11229-006-0012-2

Lin, H. (2022). Bayesian epistemology. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford encyclopedia of philosophy* (2022nd ed.). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/fall2022/entries/epistemology-bayesian/

Moretti, L., & Wright, C. (2023). Epistemic entitlement, epistemic risk and leaching. *Philosophy and Phenomenological Research, 106*(3), 566–580.

Moyal-Sharrock, D. (2004). *Understanding Wittgenstein's on certainty*. London, UK: Palgrave Macmillan.

Pettigrew, R. (2020). What is conditionalization, and why should we do it? *Philosophical Studies, 177*(11), 3427–3463.

Pritchard, D. (2015). *Epistemic Angst*. Princeton, NJ: Princeton University Press.

Quine, W. V. (1951). Two dogmas of empiricism. *Philosophical Review, 60*(1), 20–43.

Rescorla, M. (2021). On the proper formulation of conditionalization. *Synthese, 198*(3), 1935–1965.

Salvatore, N. (2015). Wittgensteinian epistemology and Cartesian skepticism. *Kriterion (Austria), 29*, 53–80. https://doi.org/10.1515/krt-2015-290205

Titelbaum, M. G. (2022a). *Fundamentals of Bayesian epistemology 1: Introducing credences*. Oxford, UK: Oxford University Press.

Titelbaum, M. G. (2022b). *Fundamentals of Bayesian epistemology 2: Arguments, challenges, alternatives*. Oxford, UK: Oxford University Press.

Volpe, G. (2021). Contrastive hinge epistemology. *Theoria, 87*(5), 1222–1249. https://doi.org/10.1111/theo.12339

Wittgenstein, L. (1969). In G. E. M. Anscombe & G. H. von Wright (Eds.), *On certainty.* Translated by D. Paul and G. E. M. Anscombe. Oxford, UK: Basil Blackwell.

Wright, C. (2004). Warrant for nothing (and foundations for free)? *Aristotelian Society Supplementary, 78*(1), 167–212. https://doi.org/10.1111/j.0309-7013.2004.00121.x