Hyperintensional Evidence and Bayesian coherence

Ted Poston
University of Alabama

December 12, 2023

Abstract

Bayesian approaches to rationality require that a person's degrees of belief be coherent. Among other implications, coherence requires that a person has the same degree of belief in every logically equivalent proposition. However, a person can have evidence for a claim without having evidence for all its propositional equivalences. This paper explores this conflict and argues that a person may be perfectly rational in virtue of responding to their evidence, even if their credences are not coherent. The paper also challenges the idea that it is always better to have more coherent credences, highlighting the fundamental role that evidence plays in rational belief.

Bayesian approaches to rationality typically begin with the requirement that a person's degrees of belief should be representable by a probability function. This is the requirement that a person's degrees of belief be *coherent*. Coherence requires, among other things, that a subject who has a specific degree of belief in the proposition p must have the same degree of belief in every proposition that is logically equivalent to p. Yet a reasonable person may have a high credence in p but a low credence in $\neg(p \rightarrow \neg(q \rightarrow q))$, even though these propositions are logically equivalent. An evidentialist may say that while the subject has evidence that p that very same evidence doesn't thereby support every logically equivalent proposition. Evidence is hyperintensional; evidence may support p but not support every proposition that is logically equivalent to p. This points to a prima facie conflict between Bayesianism and evidentialism.

Cases like the one above are typically taken to show that the Bayesian ideal of coherence is unrealistic for ordinary agents. Even so, a standard line in the Bayesian literature is that it is better to approximate the ideal of coherence. John Earman, for instance, exclaims:

> Actual inductive agents ... lack the logical and computational powers required to meet the Bayesian norms. The response that Bayesian norms should be regarded as goals toward which we should strive even if we always fall short is idle puffery unless it is specified how we can take steps to bring us closer to the goals.[1]

Julia Staffel observes that little work has been done to make good sense of this thought. Arguments for the Bayesian ideal of coherence have focused on why it is good to be coherent; not why it is good to approximate coherence. Staffel's project in *Unsettled Thoughts* is to make precise what it is to be less incoherent in service of defending the view that "Bayesian norms of rationality are relevant to human thinkers as ideals worthy of approximation."[2] This involves substantiating two claims. First, "non-ideal thinkers can be more or less irrational and that being less irrational is good."[3] Second, "we can understand the way in which being less irrational is better in terms of how far away a thinker's credences are from the ideal."[4] Her defense of these two claims is strong and comprehensive. The careful reader of *Unsettled Thoughts* is treated to a powerful case for how Bayesian norms of rationality can apply to non-ideal thinkers.

My aim in this paper is to focus attention on how even a more realistic Bayesian epistemology conflicts with evidentialism. I will be pursuing the line above that a subject can have evidence for a claim without having evidence for all its propositional equivalences. If this line is correct then it undermines a core idea in a Bayesian account of non-ideal agents, namely that agents are always irrational when their credences are not coherent. Moreover, it undermines the idea that it is always better to have credences that are more coherent. As I will explore, a person may be perfectly rational in virtue of responding to her evidence even though her credences are not coherent.

1 Intentionality, Intensionality, and Hyperintensionality

I'll be arguing that evidence is hyperintensional and that this feature of evidence makes it incompatible with Bayesian epistemology. Let's begin by reviewing the relevant terminology and making a brief case that evidence is intensional.

[1] Earman (1992, 56)
[2] Staffel (2019, 199)
[3] Staffel (2019, 4)
[4] Staffel (2019, 4)

Intentionality is a feature of minds to represent or to be about things. A belief is an intentional state; it represents the world as being a certain way. Intentionality with a 't' creates intensionality with an 's'. Consider the following example. John's neighbor is Sam Clemens. John knows a lot about Sam. He knows that Sam likes to explore the caves outside Hannibal, that he enjoys fishing from the banks of the Mississippi on a sunny afternoon, and that he meticulously maintains his white picket fence. John also enjoys reading. He has read *The Adventures of Tom Sayer* and *The Adventures of Huckleberry Finn*. He knows that Mark Twain, the author of these novels, is a famous American author. But John doesn't know that his neighbor, Sam Clemens, is Mark Twain.

John's mental file on Sam Clemens is full of information about his neighbor, but his mental file on Mark Twain is distinct containing only information gleaned from reading Twain's novels. John believes that Mark Twain is a famous author. But, even though Sam Clemens is Mark Twain, John does not believe that Sam Clemens is a famous author. The intentionality of belief creates an intensional context. The characteristic feature of an intensional context is that the truth-value of a statement is not preserved by substituting in co-referring terms. 'Sam Clemens' refers to the same person as does 'Mark Twain'. But statements regarding what John believes about Sam Clemens change their truth value when 'Sam Clemens' is replaced by 'Mark Twain'.

This example can also be used to show that evidence creates an intensional context. John has excellent evidence that Mark Twain is a famous author. He has read the books. But this is not evidence that Sam Clemens is a famous author. Evidence comes from beliefs and experiences which are intentional. It's not surprising, therefore, that evidence is intensional. Evidence is such that it is not closed under material equivalence. Moreover, this feature of evidence suggests that evidence is hyper-intensional; that is, evidence is not closed under necessary equivalence. One may have evidence that p but fail thereby to have evidence that $\neg(p \rightarrow \neg(q \rightarrow q))$. One may be told that p, but not thereby told that $\neg(p \rightarrow \neg(q \rightarrow q))$. If evidence functions in this way then it poses a significant problem for Bayesian epistemology.

The example above of truth-functionally equivalent but distinct propositions does assume that a Boolean view of propositions is false. On a Boolean view, truth-functional equivalence is an identity condition for propositions. Yet propositions function in two ways as bearers of truth-values and items that carry meaning. Differences in meaning are what explain differences in cognitive significance. A person who learns p is disposed to use p in reasoning. A person who learns p is not thereby disposed to use $\neg(p \rightarrow \neg(q \rightarrow q))$ in reasoning. A person's mental file on p is not closed under truth-functional equivalences. As I see it, what explains good rea-

soning and rational action is constrained by a person's mental representations. A strength of Staffel's approach to Bayesian norms is that it appears to be, in general, congenial to the constraints on reasoning imposed by logically limited agents. My goal is to explore to what extent this may conflict within a general commitment to evidentialism.

2 Bayesian Norms & Evidence

Staffel argues for the view that "all thinkers whose credences don't conform to the ideal norms [of Bayesianism] are irrational to some degree."[5] She explicitly recognizes that this claim conflicts with intuitive judgements about cases. A mathematician, for example, who has less than complete confidence in an unproven theorem is not irrational in any sense. Yet Staffel defends the commitment to the irrationality of incoherent credences on broader theoretical grounds.

To assess the objection that a person may be rational even if with incoherent credences, Staffel describes two cases to focus our attention.[6]

> *Uncertain Logic Student*
>
> Una is working on a homework assignment that asks her to figure out whether some set of premises P entails some conclusion C. In fact, "If P then C" is tautological, but Una has not found the proof yet. She currently assigns the conditional a middling credence of 0.5.
>
> *Certain Logic Student*
>
> Cera is working on a homework assignment that asks her to figure out whether some set of premises P entails some conclusion C. In fact, "If P then C" is tautological, but Cera has not found the proof yet. Still, she currently assigns the conditional a credence of 1.

Staffel recognizes that the intuitive response to these cases is that Una's middling credence in the tautology is rational and Cera's full credence in the tautology is irrational. But her view is committed to the claims that *Una's middling credence is, in some sense, irrational* and that *Cera's full credence is, in some sense, rational.*

She fills out this view by bringing in the distinction between propositional rationality and doxastic rationality. Propositional rationality is a matter of what the subject's evidence supports. A claim can be supported by the evidence the subject

[5] Staffel (2019, 32)
[6] Staffel (2019, 153)

possesses but the subject hasn't focused on the evidence. Doxastic rationality requires that a subject bases her credence in the appropriate way on the evidence. Given this distinction, one may hold that the Bayesian requirements of rationality are requirements of propositional rationality. Then one can understand the cases of Una and Cera as follows. Una's middling credence in the tautology 'If P then C' is propositionally irrational because it conflicts the evidence that Una possesses. Cera's full credence in the tautology is propositionally rational because, like Una, she possesses evidence that supports the tautology; but her full credence is not based on the evidence because she hasn't worked out the proof yet.[7]

Staffel makes two general claims in her reply. I quote the relevant part. Staffel writes,

> But once we see the standard Bayesian requirements as requirements of propositional rationality, they [the judgements about Una and Cera] are a lot less counterintuitive. What is propositionally rational for a thinker depends on what their evidence supports. If a claim is tautological, and its truth can thus be determined without any empirical evidence (or, alternatively, any thinker is in some sense already in possession of the needed evidence to determine its truth), then, for any thinker, assigning this claim full credence is propositionally rational.[8]

The first claim explicates propositional rationality in terms of what the evidence supports.

> (PRE) A proposition p is propositionally rational for a subject S at t with total evidence e if and only if p is supported by e.

The second claim is a substantive commitment to tautologies being supported by evidence that a subject already possesses.

> (TE) For any subject S, any tautology $\top$, and every body of total evidence e, e supports $\top$.

(PRE) and (TE) imply that Una's middling credence is propositionally irrational and Cera's full credence is propositionally rational. (TE) implies that both Una and Cera have evidence that supports the tautology. (PRE) implies that it is propositionally rational for both Una and Cera to have a credence that matches what is supported by the evidence.

[7] See Staffel (2019, 154)

[8] Staffel (2019, 154)

3 Propositional rationality and evidence

My ultimate aim is to argue that (TE) is false. But before we look at these cases, let us examine case for (PRE).

(PRE) assumes a distinction between propositional and doxastic rationality. This distinction (though using the term 'warrant' and then later 'justification') was originally introduced to make sense of judgements in which a subject holds a belief for which she has good reasons, but her belief is not based on those good reasons. Lehrer's case of the gypsy lawyer illustrates this.[9] A lawyer is hired to defend a client of a series of eight grisly murders. The evidence overwhelmingly indicates that the client is guilty of the first seven murders, and the lawyer appropriately believes that on the basis of the evidence. The lawyer, though, believes that his client is innocent of the eighth murder even though the extant evidence strongly indicates his client is guilty of that murder as well. The lawyer believes this because the cards said that his client was innocent of that crime. On the basis of this conviction, the lawyer discovers a complicated but sound path of reasoning that will exonerate his client of the eighth murder. However, given the grisly nature of the other cases, the lawyer's belief that his client is innocent of the eighth murder is sustained by the Tarot reading and not the new evidence.

Lehrer's case highlights the difference between the lawyer's psychological state of believing in his client's innocence and the evidence the lawyer possesses about his client's innocence. As described the psychological state is neither produced nor sustained by the evidence; rather it is produced and sustained by the reading of the cards. Given this difference, we distinguish between two judgements. First, the lawyer's psychological state of believing in his client's innocence isn't appropriately based on the evidence the lawyer possesses. Second, the lawyer does have good evidence that his client is innocent. The distinction between propositional and doxastic justification codifies this difference. The lawyer has propositional justification for the relevant claim but the lawyer is not doxastically justified in believing the claim.

The cases in which the distinction between propositional and doxastic justification are useful form a range of cases in which a subject possesses evidence for a claim but that evidence neither causes nor sustains the appropriate psychological state. As we will see, there are some difficult questions about whether a subject does possess evidence for a claim, but these questions do not undermine the core idea that propositional justification is tied to the evidence a subject possesses.

[9]Lehrer (1971). This distinction between propositional and doxastic justification was introduced later by Roderith Firth. See Firth (1978).

4 Evidence possession

The key question that is raised by Staffel's response is whether subjects, like Una and Cera, possess evidence supporting any tautology. It is a strong, natural judgment about the case that Una and Cera do not have evidence indicating that 'If P then C' is a tautology. Yet, often, the burden of an otherwise good theory is that it requires smoothing over such intuitions. Nonetheless, the phenomenon in the Una and Cera case reflects a wide range of cases in epistemology about what it is to possess evidence. In the following I look at three cases in which there is broad agreement about what it is to possess evidence. These cases illustrate that evidence is intensional.

4.1 Case 1: *Fingerprints*

You are the lead detective investigating the theft of Rembrant's *The Abduction of Europa* from the Getty Museum. You discover a telling fingerprint near the scene of the crime. Your forensic technician observes that the fingerprint has a radial loop near its center. She indicates this is a rare feature that will most likely uniquely pick out the culprit.

Figure 1: Fingerpint with central radial loop

As it turns out, the fingerprint is Moriarty's. You have evidence that the print has a unique radial loop. Moreover, you have as evidence that the person with that unique radial loop is the culprit. But you don't have evidence that Moriarty is the culprit, even though Moriarty is the person with that unique radial loop. What is missing is evidence that Moriarty is the person with that unique radial loop. That is a crucial, yet unknown, fact.

Fingerprints illustrates evidence can include features that pick out other states in the world, but one doesn't thereby have evidence for those other states. What is often needed is knowledge of a relevant linking fact. In this case, the relevant linking fact is that *The fingerprint is Moriarty's fingerprint.*

4.2 Case 2: *Speckled hen*

Consider the figure 2 of a speckled hen. Suppose that the hen has 57 speckles. Unless you have savant abilities, vision does not give you evidence that the hen has 57 speckles. Yet the visual experience presents the hen as having a determinate number of speckles. The specific number of speckles is beyond the power of normal visual discrimination.

Figure 2: A Speckled Hen

Speckled hen developed as a counterexample to foundationalist views of perceptual justification according to which experience of a quality thereby gave one reason that that quality was present.[10] In this case, one's visual experience of the hen represents the hen as having many speckles, but the precise number of speckles is not itself presented in experience in a way that allows one to just 'read off' from the experience the number of speckles.

AJ Ayer suggested that if one is unable to tell by visual experience itself whether a determinate quality then experience doesn't represent that determinate quality. If this is correct, then *Speckled hen* is no longer a counterexample to such foundationalist views. But Ayer's response requires that the limits of perceptual experience are the limits of perceptual discrimination. That seems dubious. Suppose you are looking for your keys. Your desk is rather cluttered. You look intently at your cluttered desk but don't see the keys. As it turns out the keys are on the desk and in plain view. You didn't notice them. As you turn away to look somewhere else, you realize that the keys are in fact in plain view. You looked at them but didn't notice them. The moment of realization is an 'aha' moment when you realize an unnoticed feature present in your experience. Ayer's view doesn't capture this mundane fact of perceptual experience.

Speckled hen fits with *Fingerprints* in that not every feature of what you have as evidence is itself evidence that you have. In *Fingerprints* you have as evidence that the culprit has a unique radial loop; it is a feature of this loop that it is Moriarty's;

[10]See Chisholm (1942).

but you do not have as evidence that Moriarty is the culprit. In *Speckled hen* you have as evidence that the hen has a determinate number of speckles; it is a feature of this evidence that the hen has 57 speckles; but you do not have as evidence that the hen has 57 speckles. *Fingerprints* fits more broadly with cases of scientific evidence where investigation and reflection are required to determine the significance of the evidence. *Speckled Hen* fits broadly with perceptual evidence where again investigation and reflection are required to determine its significance.

4.3 Case 3: *Precarious Peak*

Our last case comes from Richard Feldman's essay 'Having Evidence.'[11] Feldman describes the case.

> Suppose my friend Jones tells me that the hike up to Precarious Peak is not terribly strenuous or dangerous, that it is the sort of thing I can do without undue difficulty. Assume that Jones knows my abilities with respect to these sorts of things and that he seems to be an honest person. On the basis of his testimony, I believe that the hike is something I can do. It seems that it is rational for me to believe this proposition. But suppose I've failed to think about the time Jones told me that I could paddle my canoe down Rapid River, something he knew to be far beyond my abilities. He just gets a kick out of sending people off on grueling expeditions. If you were to say to me, "Remember when Jones lied about the canoe trip?" I'd say "Yes! How could I have failed to think of that?" Once I was reminded of this episode, it would no longer be rational for me to believe that I can complete the hike, unless I had some additional information supporting the view that Jones was not lying this time. But are the facts about the past lie part of my evidence before you remind me of them? Whether my belief is justified depends upon the answer. If this stored information is part of my evidence, then my belief is not justified, but if it is not part of my evidence, then the belief is justified.

Feldman's case illustrates a case where information is poised to be recalled. In *Precarious Peak* it is close to recollection that your friend is a practical joker about adventure trips. Unlike *Fingerprints*, there is no need to learn additional information that provides a linking fact. Unlike *Speckled Hen* the information is not beyond normal powers of discrimination. Yet it is a case that clearly illustrates that the Feldman does not possess as evidence that hike is beyond his abilities.

[11] Conee and Feldman (2004).

Of the cases considered, *Precarious Peak* is the most precarious. It may seem to some that Feldman does possess evidence that the hike is beyond his abilities. This intuition is far from universal. Moreover, a simple tweak to the case makes it evident that Feldman does not possess such evidence. Consider a variant on the case in which Feldman is placed in a MRI machine measuring brain states associated with stored memories. If Feldman were to begin to recall the relevant facts of the case, Feldman would receive a shock that would prevent him from completing that train of memories. In this case, even though the information is poised to be recalled, it cannot be recalled because of the inhibitor. The right thing to say here is that Feldman does not possess the relevant memory as evidence.

4.4 'To be determined' cases

Let's review the case of Una and Cera again. These belong to a class of cases we may call 'to be determined' cases. Staffel writes about Una's case that it "Una's case belongs to a more general class of cases in which a thinker has not yet had a chance to examine her evidence, or to reason through a problem carefully, and thus adopts a non-committal attitude in light of this."[12] Staffel's view is that these TBD cases are ones in which the subject possesses the relevant evidence but has not carefully considered it. Yet reflections on a wider range of cases in epistemology reveals that there are TBD cases in which the subject does not possess the relevant evidence. The cases of *Fingerprints*, *Speckled Hen*, and *Precarious Peak* demonstrate this.

From a Bayesian perspective, these cases are not necessarily troublesome. It's open for the Bayesian to hold a theory of propositional content that permits learning that the fingerprints have a unique radial loop without thereby learning that Moriarty's fingerprints have this unique radial loop. Similarly, a Bayesian may adopt a view of perceptual experience and perceptual content that allows one to learn a more general fact about a perceptual experience without learning every fact about the perceptual presentation. And similarly, a Bayesian can adopt a view of memory states that not every stored belief counts as evidence possessed.

What is more troubling here is that the Bayesian view of the a priori is different from the above. The coherence norm is a constraint on rationality such that any departure from coherence is a departure from full rationality. This is why intensionality per se is not particularly troubling for Bayesians; it's open to deal with those problems from a theory of content. Yet hyperintensionality is troubling because of the way it interacts with the norm of coherence. If evidence is hyperintensional then one can have evidence for p and not have evidence for a propositional equivalence of

[12]Staffel (2019, 155)

p. I've given a case of this at the start of the paper. A person may gain evidence that p without thereby gaining evidence that $\neg(p \rightarrow \neg(q \rightarrow q))$. This case differs only from the wider range of cases in being a priori. Yet in this case and the above cases, they require successfully completing a deliberative process to come to possess the relevant evidence. Accordingly, there is no relevant epistemological difference between the intensionality of evidence and the hyperintensionality of evidence.

5 Evidential support, entailment, and Bayesian methods

I've been pursuing the idea that hyperintensional evidence makes trouble for a core feature of Bayesian accounts of rationality. In the following, I discuss a broader set of issues relating to hyperintensional evidence, evidential support, and how we should conceive of Bayesian methods. Staffel's discussion of different ways to think of Bayesian methods in her second chapter serves as a useful tool for thinking about hyperintensional evidence.

Given some evidence, it's natural to wonder what that evidence supports. Suppose evidence, e, implies p. It's instinctual to think that e supports p. Moreover, it's intuitive to think that entailment is the limiting case of evidential support. One might, thus, be attracted to the following principle of Evidential Support Closure.

> (ESC) For any subject S and every body of total evidence e, if e entails p then e supports p.[13]

(ESC) fits in with the guiding intuition between Hempel's Equivalence Condition.[14] Hempel held that (dis)confirmation is closed under logical equivalence. My sense of Hempel's rationale for the equivalence condition was to make sense of the role of logic in scientific experiment. A theory is tested by deriving a potential observational consequence from it. It does not matter for the purpose of the logic of experiment whether one worked with the theory as formulated or a logically equivalent formulation of the theory.

How does hyperintensional evidence fit within this broader view of (ESC) and Hempel's Equivalence Condition? The challenge here is that by focusing on the way evidence is used to gain further confirmation or disconfirmation, we are driven back to think of evidential support as closed under entailment. The problem with this

[13](TE) is a special case of (ESC).

[14]Thanks to an anonymous reviewer for this point.

suggestion is that the logical entailments that are used in both testing theories and expanding evidence are the ones which are known. In a scientific context, one uses a model in which the relevant logical contributions are already known. Thus the relevant principle is that evidential support is closed under known implication.

> (ESKC) For any subject S and every body of total evidence e, if S *knows* that e entails p then e supports p for S.

This fits more generally with the use of Bayesian models in science. Bayesian models are superb in teasing apart subtle probabilistic relationships. Take, for instance, the standard Bayesian solution to the ravens paradox.[15] The paradox focuses on the unintuitive result that the observation of a white shoe confirms that all ravens are black. This result follows from two natural principles on confirmation.

> Nicod's Condition: Universal laws are confirmed by instances.

Nicod's condition, in particular, states that a law of the form '$(\forall x)(Fx \rightarrow Gx)$' is confirmed by an instance 'Fa & Ga'.

The next principle is Hempel's Equivalence Condition.

> Equivalence Condition: For any three propositions e, h_1, h_2, if e confirms h_1 and h_1 is equivalent to h_2 then e confirms h_2.

The paradoxical result follows from these two conditions. All ravens are black is equivalent to all non-black things are non-ravens. A white shoe is in instance of a non-black, non-raven, and so confirms that all non-black things are non-ravens. Thus, by the equivalence condition, a white shoe confirms that all ravens are black.

The standard Bayesian solution to the paradox begins by observing that differences in confirmation are affected by what is in a person's background knowledge. If, for example, one already knows that an object to be sampled is not a raven then observing its color is irrelevant to whether all ravens are black. Moreover, IJ Good showed that there are some sets of background information on which observing a black raven disconfirms that all ravens are black.[16] The key to the Bayesian solution is to observe that given that (1) there are many more non-black objects in the universe than ravens and (2) that the probability of observing a non-black object or a

[15] For an excellent overview and a novel contribution to the Bayesian view, see Fitelson and Hawthorne (2010).

[16] See Good (1966)

raven is conditionally independent from all ravens are black, it follows that observing a black raven more strongly confirms that all ravens are black than observing a white shoe.[17] Thus, sensitivity to background knowledge ameliorates the paradox of confirmation. But it takes the fine-grain of a Bayesian approach to see the subtle differences in confirmation.

I conceive Bayesianism as a useful tool to focus on subtle probabilistic dynamics in an otherwise well understood theory. For the reasons I have been offering, I do not think that Bayesianism is an adequate a theory of rational credences for ordinary agents. Staffel helpfully highlights several ways of conceiving of Bayesian models in chapter two of her book. She identifies three views: the ideal agent view, the scientific idealization view, and the systematization view. An ideal agent view thinks of Bayesian norms as norms for ideally rational agents. A scientific idealization view conceives of Bayesian norms as constructing false models which may be useful for various purposes. The ideal agent view doesn't have a clear normative application to ordinary agents, and the scientific idealization view isn't in the business of issuing normative judgements about ordinary agents.

Staffel argues that the Bayesian model is best understood as a systematization view. She describes it as follows.

> The norms are derived by thinking about the characteristic role our credences are supposed to play in our thinking and decision-making. Once we have a systematic understanding of the role of our credences, we can develop norms that our credences should obey in order to perform their roles in the best possible ways. In developing the norms, we abstract away from limiting factors that interfere with credences playing this role perfectly, such as processing or time limitations or possibilities of error ... On this view, we can also make sense of the idea that non-ideal thinkers are better off the more closely they approximate rational ideals. If fully rational credences are best at performing their functional roles, it is at least prima facie plausible to assume that the less rational one's credences are, the worse they are at performing these roles, and the more rational they are, the better they are at performing them.[18]

I have some considerable sympathy with the systematization view of Bayesian models. The key question is how this view may fit with broader reflections about the nature of evidence. For the reasons I've been discussing, a person may have evidence

[17] For details see Fitelson and Hawthorne (2010, 255–258).

[18] Staffel (2019, 24–5)

for a claim without having evidence for a truth-functionally equivalent claim. This doesn't seem to be a limiting case that one can abstract away from for the purposes of a normative theory of rationality. Rather it seems to be a feature of evidence for ordinary agents. There may be some space between a scientific idealization view and a systematization view that allows for some useful falsehoods about agents and still deliver robust norms. If so, then there may be room in a Bayesian mansion for hyperintensional evidence.

6 Conclusion

Staffel's book *Unsettled Thoughts* is an excellent contribution to the literature on Bayesian epistemology. Her book is the first comprehensive development of the Bayesian thought that it is better to be less incoherent. Moreover, she has put her finger on the important and under-explored phenomenon of 'to be determined' cases. I've suggested an evidentialist response to these cases. If this line is correct, then the TBD cases reveal a deeper conflict between Bayesianism and evidentialism.

References

Chisholm, R. (1942). The problem of the speckled hen, *Mind* **51**: 368–373.

Conee, E. and Feldman, R. (2004). *Evidentialism: Essays in epistemology*, Oxford University Press.

Earman, J. (1992). *Bayes or Bust?*, MIT Press.

Firth, R. (1978). Are epistemic concepts reducible to ethical concepts?, *in* A. Goldman and J. Kim (eds), *Values and Morals: Essays in Honor of William Frankena, Charles Stevenson, and Richard Brandt*, Dordrecht: Kluwer Academic Publishers, pp. 215–229.

Fitelson, B. and Hawthorne, J. (2010). How bayesian confirmation theory handles the paradox of the ravens, *in* E. Eells and J. Fetzer (eds), *The Place of Probability in Science*, Vol. 284, Springer, pp. 247–275.

Good, I. (1966). The white shoe is a red herring, *British Journal for the Philosophy of Science* **17**(4): 1966.

Lehrer, K. (1971). How reasons give us knowledge, or the case of the gypsy lawyer, *The Journal of Philosophy* **68**(10): 311–313.

Staffel, J. (2019). *Unsettled Thoughts: A Theory of Degrees of Rationality*, Oxford University Press.