Inferential Evidence

Jeff Dunn

Forthcoming in *American Philosophical Quarterly*, please cite published version.

1 Introduction

Consider:

> **The Evidence Question:** When, and under what conditions does an agent have proposition E as evidence (at t)?[i]

Timothy Williamson's ([2000]) answer to this question is the well-known E=K thesis:

> **E=K:** E is a member of S's evidence set at t iff S knows E at t.

I will argue that this answer is inconsistent with the version of Bayesianism that Williamson advocates. This is because E=K allows an agent to garner evidence via inductive inference whereas standard Bayesian views disallow such a thing. Since Williamson's version of Bayesianism shares the key features with the standard Bayesian view, there is an inconsistency.

Others have considered problems that arise for E=K. Comesaña & Kantin ([2010]) argue that E=K runs into problems with Gettier cases and a closure principle for justification. Their arguments, however, criticize E=K because of the left-to-right implication that requires that all evidence is known. My criticism, in contrast, is that E=K is mistaken because of the right-to-left implication that requires that all known propositions are evidence. Dodd ([2007], pp. 644-9) and Littlejohn ([2011], pp. 247-8) both take aim at the right-to-left implication of E=K.

Both do so by considering cases of inferential evidence and inferential knowledge, and thus their arguments are, in a way, similar to mine. However, their arguments do not address the problems that arise for E=K because of a commitment to a Bayesian picture of belief change, and so in this way are considerably different.

Further, because the argument in this paper focuses on a conflict between Bayesianism and inductive inferential evidence, it is of interest independent of E=K. Quite apart from E=K, it shows that one cannot maintain two attractive theses: that agents can sometimes garner evidence via inductive inference, and that the standard Bayesian view provides a model of rational belief and belief change. Though it is common for Bayesians to think of evidence propositions as coming from direct perception (along the lines of Carnapian protocol sentences), there is some intuitive force to the idea that evidence can come via inductive inference. When light filters through the curtains in the morning, I do not directly perceive that the sun has risen. But if someone were to ask me what evidence I have for thinking it is morning, it is natural to say my evidence is that the sun has risen. So, the forced choice between Bayesianism and inductive inferential evidence is interesting in itself. In the last sections of this paper, I show how one can modify the Bayesian view to allow evidence that is garnered via inductive inference, and I consider whether these modifications are amenable to Williamson.

2 E=K and Inferential Evidence

We know some propositions because we infer them from other propositions. Sometimes the inference is deductive—I know that there is someone in the room as a result of inferring it from the known fact that you are in the room. Sometimes the inference is inductive—I know that the coming winter will be colder than the summer as a result of inferring it from known facts about the past. Since E=K says that all known propositions are evidence, it has the consequence that a proposition can come to be evidence as the result of an inference.

One might balk at this quick argument. Suppose that it is true that every statue of Napoleon is physical. It may be true that this object is a statue of Napoleon as a result of careful planning, but it needn't be true that this object is thereby physical as a result of careful planning. Similarly: suppose that E=K is true. One might argue that just because P is known as a result of inductive inference, it needn't be the case that P is evidence as a result of inductive inference. In fact, however, this objection fails. According to E=K it is a proposition's status as knowledge that makes it evidence. E=K is part of Williamson's "knowledge first" epistemology. If P's status as knowledge brings about P's status as evidence, then it seems innocuous to say that any route to knowledge is thereby a route to evidence.

3 Bayesianism and Inferential Evidence

3.1 Bayesianism

The standard Bayesian approach to epistemology is incompatible with this kind of inferential evidence. To see this, we must briefly examine two of its key features. Fundamental to the Bayesian approach is the claim that beliefs come in degrees and that, to be rational, these degrees of belief (henceforth: 'credences') must be representable by a probability function. Second, the Bayesian approach gives an account of how beliefs should change over time. The standard view is that a rational agent updates her credence function upon the receipt of evidence. Conditionalization tells us how this goes:

Conditionalization: $cr_{t1}(\cdot) = cr_{t0}(\cdot|E_{t1})$

In this expression E_{t1} is all the new evidence the agent learns at t1 and $cr(A|E) =_{df} cr(A \wedge E)/cr(E)$. An important implication of Conditionalization is that the proper response to one's evidence is to give it full credence (because $cr(E|E) = 1$). In what follows, I'll use the term 'Standard Bayesianism' to refer to any view with these two features.[ii]

3.2 Inference and Standard Bayesianism

To show a conflict between Standard Bayesianism and inferential evidence, something must be said about inference in the Bayesian framework. On one natural way of interpreting the Bayesian formalism, a credence in a proposition can change in two ways:

1. because the proposition becomes evidence,
2. because of a Conditionalization update.

For example, suppose that at t0 I am such that:

t0: $cr(A|E) = 0.9$ $\quad$ $cr(A) = 0.5$ $\quad$ $cr(E) = 0.5$

Perhaps E is the proposition that there is white smoke coming from the Sistine Chapel, and A is the proposition that a new Pope has been elected. Suppose that at t1 I look above the Sistine Chapel and see white smoke, thus acquiring E as evidence.

My credence in E goes to 1, and as a result of Conditionalization, my credence in A goes to 0.9. At t1, then, things look like this:

t1: $cr(A|E) = 0.9$ $cr(A) = 0.9$ $cr(E) = 1$

In this simple case, two credences have changed: my credence in E and my credence in A. The change in A is most like an inference, since it is a change in a credence, brought about solely in virtue of my other credences (namely, my credence in E). Note that this is how inference works for all-or-nothing beliefs. If A is inferred from E then one's belief that E results from one's belief in A (and perhaps also one's belief about the relationship between E and A).

One artificial feature of the Bayesian framework is that the changes in credence to E and to A are represented as taking place instantaneously. Nevertheless, it is natural to conceptually distinguish these kinds of change. The simple example above is exactly how a Bayesian would model what it is to inductively infer that a new Pope has been elected on the basis of the evidence that white smoke is coming from the Sistine Chapel. My rough proposal, then, is to say that within Bayesian models a shift in opinion is the result of an inference when it is a shift in opinion similar to the one with respect to A in the above example.

More carefully, we can distinguish two ways that opinions shift during a Conditionalization update. First, a shift in opinion can itself reflect the acquisition of evidence. This is what happens when the credence in E goes from 0.5 to 1. Second, a shift in opinion can be brought about by the evidence. This is what happens when the credence in A goes from 0.5 to 0.9 in virtue of the fact that $cr_{t1}(E) = 1$ and $cr_{t0}(A|E) = 0.9$. This gives us at least a necessary condition for a change in credence to be the result of an inference: one's credence in A at t1 is the result of an inference from one's full credence in E at t1 only if $cr_{t1}(A) = cr_{t0}(A|E)$.

Three points should be made clear. First, some might think that the term 'inference' is inappropriate here. One might think that an inference can obtain only between *believed* propositions. For example, one comes to believe Q via inference from one's belief in P and one's belief in $P \rightarrow Q$. The Bayesian framework doesn't have (full) beliefs in this sense, and so perhaps it's not quite right to say that the Bayesian framework countenances inferences, *per se*. I can grant this point. What is important for my purposes is that shifts in opinion that are brought about by evidence are very similar to inferences, and in fact how a Bayesian would model intuitive cases of coming to believe something via inference rather than, say, via direct observation.

Second, it should be noted that one doesn't have to interpret the Bayesian formalism in just the way that I have suggested. In the change from t0 to t1 I proposed to distinguish between the way in which the credence assigned to E changed and the way in which the credence assigned to A changed. But one needn't do that. For instance, one could say that my credence in E at t1 is the result of the fact that E is evidence and that $cr_{t0}(E|E) = 1$. Similarly, my credence in A at t1 is the result of the fact that E is evidence and that $cr_{t0}(A|E) = 0.9$. On this interpretation, the changes in credence to A and to E are seen as very similar. My argument will go through on this interpretation, so long as one grants that the change in credence to A is the result of an inference from E.[iii] However, I think that this is an unnatural interpretation of the formalism. It is unnatural because we lose the distinction between changes in credence that are brought about by other credences and changes in credence that are brought about more directly. This seems like a distinction worth preserving. Consider the Sistine Chapel case. The way in which I come to my opinion that there is white smoke is very different from the way in which I come to my opinion that a new Pope has been elected. The latter opinion seems to be dependent in an important way on the former opinion. It is this very dependence that makes the latter change seem to result from an inference and the former change to not result from an inference. So, I'll continue to mark this distinction.

There is one final, important point. In the previous section I noted that a distinction can be drawn between two kinds of inference. First, a person can come to believe a proposition on the basis of what she views as an inductive inference. This is a situation where the agent sees the inference as less than fully certain. Second, a person can come to believe a proposition on the basis of what she views as a deductive inference. This corresponds to an inference the agent thinks of as fully certain.[iv] On the interpretation of the Bayesian formalism that I have suggested, we can draw this same distinction. An inductive inference from E to A corresponds to a situation where $cr_{t1}(E) = 1$ and $cr_{t1}(A) = cr_{t0}(A|E) < 1$. A deductive inference corresponds to the same scenario but where $cr_{t1}(A) = cr_{t0}(A|E) = 1$.

4 The Argument

I will now show the problem that arises if one thinks that a proposition like A is evidence in virtue of an inductive inference.

As above, suppose that E is evidence and so assigned full credence. Suppose for reductio that A is evidence for me and also that it is evidence in virtue of being inductively inferred from E. My credence in A is to be regarded as resulting from an inductive inference from my credence in E. Thus, given the necessary condition on a change in credence being an inference, the result of this inference is that $cr_{t1}(A) = cr_{t0}(A|E)$. Since the inference to A is meant to be inductive, it must be that $cr_{t0}(A|E) = b \neq 1$. If A is evidence, however, it follows from Conditionalization that $cr_{t1}(A) = 1$. But this contradicts that $cr_{t1}(A) = b \neq 1$. On the other hand, we could suppose that $cr_{t1}(A) = 1$, but if we do so then the credence in A isn't the result of an inductive inference from E. But it is just that inference in virtue of which A is supposed to be evidence. So, A cannot be both inductively inferred and evidence at t1.

Notice that there is no such problem if the inference is seen as deductive. In that case $cr_{t0}(A|E) = b = 1$. If A is inferred from my credence in E, then $cr_{t1}(A) = 1$, and since A is evidence Conditionalization says that $cr_{t1}(A) = 1$. There is no conflict here between a proposition being evidence and one's credence in that proposition being the result of a deductive inference.

The argument, then, shows that there can be no evidence via inductive inference on the Standard Bayesian view, though it does not tell against evidence being obtained via deductive inference. Recall, however, that E=K is committed to there being evidence that results from inductive inference, since it is possible to come to know a proposition via inductive inference. In particular, I could come to know that a new Pope had been elected by inferring it from the white smoke I see rise from the Sistine Chapel. Thus, the argument here shows that E=K is inconsistent with this kind of Standard Bayesian view. In the next section I will argue that Williamson is committed to something very similar to this view, and so Williamson's commitments are inconsistent with his E=K thesis.

Before getting to that, I will consider two responses to this argument. The first response seeks to show how inductive inferential evidence is in fact consistent with the Standard Bayesian view. According to this response we rethink how the inference from E to A is to be modeled. On this alternative view there are three stages to consider. Initially, as before, I am such that:

t0: cr(A|E) = 0.9 cr(A) = 0.5 cr(E) = 0.5

Then, at t1, I get E as evidence and conditionalize:

t1: cr(A|E) = 0.9 cr(A) = 0.9 cr(E) = 1

Now, however, A is evidence, and so at t2 things are as follows:

t2: cr(A|E) = 1 cr(A) = 1 cr(E) = 1

The idea is that the change from t0 to t1 shows that the change in my credence in A was the result of an inductive inference from my credence in E, and the change from t1 to t2 shows that my credence in A really is evidence since it is assigned credence 1.[v]

The problem with this response is that the change from t1 to t2 is now completely mysterious. According to this view, the initial inference from E gets me to 0.9 credence in A, and then, subsequently, with no outside input, this credence gets bumped to 1. I see how one can say that a change in the credence to A from 0.5 to 0.9 is the result of an inference from E. But I do not see how one can say that this extra bump from 0.9 to 1 is in any way the result of an inference from E. By t2 A may be evidence, but not in virtue of any kind of inference from E. Thus, this response fails.

The second response is to question my interpretation of the Bayesian formalism, in particular my account of inference. One could maintain that the Bayesian picture is one that completely eschews inference. On this view, there are shifts in credence from one time to the next, and these changes conform to Conditionalization, but that is all we can say. This kind of view might be attractive to those who want to use the Bayesian formalism only as part of a theory of confirmation that tells us which propositions confirm which others. If there are no inferences on a Bayesian view, then one might think the above problem fails to arise.

But this is mistaken. If we want to integrate the Bayesian view with a view like Williamson's E=K, then the problem is still there even if we don't call any changes in credence 'inferences'. Suppose that I see the white smoke come out of the Sistine Chapel and as a result come to know that a new Pope has been elected. How did I come to know that? On the view under consideration this must have something to do with the fact that the proposition that there is white smoke strongly confirms for me the proposition that a new Pope has been elected. E=K further says that the proposition that a new Pope has been elected is evidence for me. It is now something on which I should conditionalize. But then we have the same conflict: propositions can be known because they are highly, though not decisively, confirmed by other things we know. But there are highly, though not decisively, confirmed propositions that cannot be evidence.

5 Williamson and Bayesianism

The argument above shows that within the confines of Standard Bayesianism, a proposition cannot be evidence solely in virtue of being inductively inferred from some other propositions. What features of Standard Bayesianism generate this conflict? We get a conflict so long as the view in question is committed to the following:

1. If an agent has some evidence, then this evidence rationally motivates changes in her doxastic state, in accordance with how relevant the evidence is to other propositions. (In the argument, this follows from commitment to Conditionalization.)
2. The proper response to evidence you have is to fully believe it. (This follows from commitment to Conditionalization, too. In the argument, this allows one to say that inductively inferred evidence will not be treated as evidence, and if it is treated as evidence, then it is not inductively inferred.)
3. If a doxastic state with content A is inferred from another doxastic state with content E then the doxastic state with content A changes as a result of the doxastic state with content E, and proportional to how relevant E's truth is to A's truth. (In the argument, this is the claim that one's credence in A at t1 is the result of an inference from one's full credence in E at t1 only if $cr_{t1}(A) = cr_{t0}(A|E)$.)

Williamson ([2000]) is committed to these features. Instead of speaking of rational credences, he speaks of evidential probabilities. Though he rejects Conditionalization, he offers in its place what he calls 'ECOND' ([2000], p. 220).

ECOND is extremely similar to Conditionalization. There are two key differences. First, ECOND allows an agent to lose evidence, while Conditionalization does not. This has no effect on the argument. Second, ECOND is defined for evidential probabilities rather than credences. Evidential probabilities reflect the true evidential support relations between propositions rather than the agent's own subjective view about the relevance of propositions to each other. However, both Conditionalization and ECOND maintain that an agent should update her credences/evidential probabilities on evidence received in accordance with her conditional credences/conditional evidential probabilities. The fact that credences are subjective and evidential probabilities are objective does not alter the fact that both views say an agent must update her doxastic state on evidence received in accordance with how relevant that evidence is to other things believed. The only difference is how 'relevance' is understood. So, Williamson is committed to 1.

ECOND is also similar to Conditionalization in that it requires evidence to receive full credence/evidential probability. Williamson is clear that he rejects evidence that comes in degrees ([2000] pp. 213-221). Thus, Williamson also accepts 2.

Consider finally point 3. Williamson could perhaps deny my proposal concerning how to understand inference in a Bayesian framework. According to this proposal, recall, one's credence in A at t1 is the result of an inference from one's full credence in E at t1 only if $cr_{t1}(A) = cr_{t0}(A|E)$.

One way to challenge this claim is to deny that there are any inferences in a Bayesian framework. At the end of Section 4 I explained why this doesn't work. A different way to challenge this claim is by providing some other way of modeling inference within a Bayesian framework. However, I think this would be a challenging thing to do. For instance, suppose that one says that a change in a credence towards a proposition is the result of an inference only if that change is caused by some other credence(s). This picture would allow for a change in a credence to be an inference, even if it were not in accordance with Conditionalization. However, this picture does not distinguish genuine inference from mere causation. For certainly a belief can cause a change in another belief without this change being an inference. Appealing to Conditionalization (or ECOND) makes the change dependent on the conditional credences (or evidential probabilities) the agent has and so is one way to make clear that the content of the propositions believed are important to distinguishing inferential relations between doxastic states from mere causal relations.

Thus, pending an alternative proposal about how to understand inference in a Bayesian framework, the view Williamson presents in his ([2000]) appears to be inconsistent. E=K does not fit with the Standard Bayesian framework that Williamson advocates.[vi] The argument is of wider interest than this, however. For many might feel attracted to the Standard Bayesian framework, and also feel that there is nothing wrong with inductive inferential evidence. The argument above forecloses the possibility of combining these two plausible ideas.

6 Responses

There are several ways that one can respond to the conflict between Standard Bayesianism and inductive inferential evidence. One way of approaching this is to notice that the argument shows that one cannot accept both of the following:

> **Certain Evidence:** evidence propositions are assigned full credence.
>
> **Inductive Inferential Evidence:** some evidence can be garnered via inductive inference.

Since Williamson ([2000]) accepts both, that view is in trouble. But this leaves an important question: what is the best fix for a Williamsonian view?

6.1 Rejecting Inductive Inferential Evidence

The first option to consider is the rejection of evidence garnered by inductive inference. This would require a modification of E=K, since as argued in Section 2, knowledge can be acquired via inductive inference. But before considering such a modification, we should consider whether it is plausible to reject inductive inferential evidence. Alexander Bird ([2004]) presents a scenario designed to show that it is not.

> **Poisonous Liquid**
>
> Suppose that there are many tests one could do to determine whether or not a certain liquid is poisonous. One test is to dip a piece of reactive paper into the liquid and observe what color the paper turns. Suppose I choose to do this particular test. Grant that I get as evidence the proposition that the test paper turned blue. From this I infer P: the liquid is poisonous. Suppose that I come back to the lab a week later. I remember that the liquid is poisonous,

> but I have forgotten which test I did. That the liquid is poisonous cannot be evidence for me if evidence is non-inferential. But this is implausible, for if that proposition is not evidence, then I have no evidence for the proposition that I should not drink the liquid. But I do.

Bird maintains that we must allow that P is evidence in this situation, even though it is inductively inferred. This is because if we deny that P is evidence, we are forced to say that I have no evidence for the proposition that I should not drink the liquid. If Bird is right, then rejecting Inductive Inferential Evidence is not a viable solution to the problem raised in this paper.

There are, however, things that a defender of the Standard Bayesian framework can say about Bird's case without countenancing P as evidence. One simple response is to maintain that a week after doing the test P *is* non-inferential evidence. Although P was initially inferred from propositions about the reactive paper, when P is stored in memory and then recalled, it may no longer be correct to think of P as inferential. Conee & Feldman ([2008]) analyze memorial evidence in a way that is similar to this.

A different response is to adopt a close relative of a distinction made by Conee & Feldman ([2008]). They distinguish between *ultimate evidence* and *intermediate evidence*. As they use these terms, ultimate evidence is the fundamental evidence an agent has. In their view, ultimate evidence is always experiential: my feeling of warmth, for instance, is ultimate evidence (perhaps for the proposition that it is warm). Intermediate evidence, on the other hand, is the kind of evidence that one has when one believed proposition is evidence for another believed proposition: the proposition that it is warm is intermediate evidence (perhaps for the proposition that it is summer). Standard Bayesianism probably doesn't fit well with Conee & Feldman's insistence that ultimate evidence is always experiential. Evidence, on a Standard Bayesian view is propositional (see endnote *i*). However, the Standard Bayesian can mark a similar distinction. If E is *intermediate evidence* for H this just means that E provides some support for H. To say that E is *ultimate evidence*, on the other hand is to say that one directly received the information that proposition E expresses. Indeed, within the Bayesian framework there is a particularly elegant description of how these two notions can coexist. E is ultimate evidence just in case one should conditionalize on E, and E is intermediate evidence for H just in case $cr(H|E) > cr(H)$.[vii]

With this distinction in hand, the Bayesian can maintain that P is *intermediate evidence*, which supports the fact that one should not drink the liquid. It is also true that I am reasonably confident that P is true. This is all that is needed to explain how I can use P to become reasonably confident that I should not drink the liquid. This is consistent, however, with P not being *ultimate evidence*. Now, one might demand a story about how I can be reasonably confident in P if P isn't ultimate evidence. In the story Bird gives, P cannot be supported by propositions about the tests I performed, since, by hypothesis, I have forgotten those. But P could be supported by some other propositions that I have as ultimate evidence, for instance, by propositions about what I seem to remember, and general propositions about the reliability of my memory.[viii]

So, we have two kinds of response to Bird's argument. The first response maintains that propositions like P really are ultimate evidence. The second response maintains that P isn't ultimate evidence but disarms the intuition that it is a kind of evidence by noting that it is intermediate evidence. On either way of going, however, the Bayesian maintains that there is no inductive inferential *ultimate* evidence.[ix]

So, is this a way out for Williamson? On its own it is not. E=K is supposed to tell us which propositions are evidence, in the sense that we should conditionalize on them. So, E=K must be telling us about *ultimate evidence*. It says that every known proposition is ultimate evidence. But that's mistaken. As the main argument of this paper shows, propositions that are known via inductive inference are not propositions that one ought to conditionalize on. So, this strategy alone does not offer a response for Williamson.

What is needed is a modification to E=K. Here is a simple fix:

E = non-inferential K: E is a member of S's (ultimate) evidence set at t iff S knows E at t and E is not known by S at t solely via inductive inference.

This view of evidence is similar to a view that has been defended by Clayton Littlejohn ([2011]). One way of seeing my argument, then, is as providing an avenue of support for that view. The view does come with some costs, however. It lacks the elegance and simplicity of E=K. It also conflicts with some of the arguments Williamson offers in his ([2000])—in particular those that maintain that all knowledge is evidence. It does, however, avoid the argument raised in this paper and in this section we've seen how a Bayesian can respond to challenges that one *must* allow non-inferential evidence.

6.2 Rejecting Certain Evidence: Graded Inferential Evidence

The other way out of the argument is to accept ultimate evidence that is acquired via inductive inference and reject Certain Evidence. This is to allow that evidence need not receive full credence. Richard Jeffrey ([1965]) famously presents a generalization of Conditionalization that models graded evidence. The updating rule is:

Jeffrey Conditionalization: $cr_{t1}(\cdot) = \sum_i [cr_{t0}(\cdot \mid E_i) \times cr_{t1}(E_i)]$

In this expression, the E_i must form a partition. It is this partition that represents the agent's evidence at t1. With Jeffrey Conditionalization an agent's evidence is not a single proposition or conjunction of propositions. Perhaps the best way to think of evidence on this picture is as non-propositional. Jeffrey's preferred interpretation is that the evidential impact of an experience is represented by a shift in credence over a partition that is assigned weights summing to 1. For instance, suppose that for a given experience the evidence partition is {E,¬E}. The evidential impact of this experience might correspond to E being assigned a credence of 0.7 and ¬E being assigned a credence of 0.3. One way to put this is to say that the agent's evidence is {⟨E, 0.7⟩, ⟨¬E, 0.3⟩}. On this picture, at each moment an agent has a certain evidence partition with weights assigned to the members of that partition. This weighted partition represents the evidential impact on the agent at a time. Although Jeffrey thinks of this evidential impact as coming solely from experience, the advocate of inductive inferential evidence could maintain that this evidential impact can come in part via inductive inferences.

Consider again the Sistine Chapel example. I see white smoke rising. Suppose that in virtue of this visual experience, the evidential impact is $\{\langle E,1\rangle, \langle \neg E,0\rangle\}$. That does not exhaust the story about my evidence, however. I also infer from my new credence in E that a Pope has been elected (A). This inference doesn't result in A being given full credence. Rather, the inference results in A being assigned some credence less than 1. But this inference itself is part of the evidential impact on me at that time. At the end of all of this, then, my evidence partition is: $\{\langle (E \wedge A), n\rangle, \langle (E \wedge \neg A), m\rangle, \langle (\neg E \wedge A), r\rangle, \langle (\neg E \wedge \neg A), s\rangle\}$. We know that E receives full credence, so that $n + m = 1$. The evidential impact with respect to A was due solely to an inference from E, so we can work out the credence that A should be given according to the equation:

$$cr_{new}(A) = cr_{new}(E) \times cr_{old}(A|E) + cr_{new}(\neg E) \times cr_{old}(A|\neg E)$$

Supposing that $cr_{old}(A|E) = 0.9$, we can work out that $n = 0.9$, $m = 0.1$, $r = s = 0$.[x]

Notice that on this view the change in credence with respect to A/¬A does seem to be the result of an inductive inference (from E). There is no extra boost that my credence in A gets outside of the impact that E has. Further, A/¬A can be part of the agent's evidence partition since there is no requirement that evidence be assigned full credence. Thus, this looks like a way of countenancing *ultimate* evidence that is a result of inductive inference: I update on a partition that includes A/¬A. It allows this by allowing graded evidence.

There are important questions about this model and about Jeffrey Conditionalization in general.[xi] However, without getting into the technical details we can see that this approach is not going to be one that Williamson can adopt.

Williamson ([2000] pp. 213-221) is clear that graded evidence is not something to be endorsed. But even if he were to change his mind on this, E=K does not sit well with graded evidence. The things that we know are propositions. E=K identifies these things with our evidence. So, E=K says that our evidence is propositional. But graded evidence isn't propositional: it is a partition with weights over a space of propositions. So, if E=K is true, then evidence isn't graded. In the previous section I noted that there is a simple modification to E=K that allows Williamson to reject Inductive Inferential Evidence. There is, however, no simple modification to E=K that allows the rejection of Certain Evidence.

7 Conclusion

In this paper I have considered the acceptability of inductive inferential evidence, arguing that the central notion of evidence within the Standard Bayesian framework is one that prohibits such evidence. This shows something interesting and surprising: one cannot consistently combine Standard Bayesianism with the claim that evidence can sometimes be acquired via inductive inference. I argued that this reveals an inconsistency in the view presented by Williamson in his ([2000]). E=K does not sit well within a Standard Bayesian framework. I've also considered two different ways one could respond to the conflict between Standard Bayesianism and inferential evidence. The first response rejects (Inductive Inferential Evidence) and uses the distinction between ultimate and intermediate evidence to explain away the apparent problems with this rejection. I argued that a modified version of E=K allows one to maintain a view similar to the one in Williamson ([2000]). The second, more radical, response rejects (Certain Evidence). This response allows for ultimate evidence that is garnered via inductive inference, but it moves one further away from Standard Bayesianism and I argued that it is not the kind of response that is amenable to E=K.

There is, of course, a large unanswered question in the background of all this: is there ultimate evidence that is garnered via inductive inference? I haven't taken a stand on this. What I have shown is that if there is, then Standard Bayesianism is inadequate.

References

Bird, A. [2004]: 'Is Evidence Non-Inferential?', *The Philosophical Quarterly*, 54, pp. 252-265.

Comesaña, J. & Kantin, H. [2010]: 'Is Evidence Knowledge?', *Philosophy and Phenomenological Research*, 80, pp. 447-454.

Conee, E. & Feldman, R. [2008]: 'Evidence', in Q. Smith (ed.) *Epistemology: New Essays*, Oxford: Oxford University Press.

Christensen, D. [1994]: 'Conservatism in Epistemology', *Nous*, 28, pp. 69-89.

Dodd, D. [2007]: 'Why Williamson Should Be A Sceptic', *The Philosophical Quarterly*, 57, pp. 635-649.

Doring, F. [1999]: 'Why Bayesian Psychology is Incomplete', *Philosophy of Science*, 66, pp. S379-S389.

Field, H. [1978]: 'A Note On Jeffrey Conditionalization', *Philosophy of Science*, 45, pp. 361-367.

Jeffrey, R. [1965]: *The Logic of Decision*, New York: McGraw-Hill.

Klein, P. [1981]: *Certainty: A Refutation of Scepticism*, Minneapolis, MN: University of Minnesota Press.

Lange, M. [2000]: 'Is Jeffrey Conditionalization Defective by Virtue of Being Non-Commutative? Remarks on the Sameness of Sensory Experience', *Synthese*, 123, pp. 393-403.

Littlejohn, C. [2011]: 'Evidence and Knowledge', *Erkenntnis*, 74, pp. 241-262.

Wagner, C. [2010]: 'Jeffrey Conditioning and External Bayesianity', *Logic Journal of the IGPL*, 18, pp. 336-345.

Wagner, C. [2002]: 'Probability Kinematics and Commutativity', *Philosophy of Science*, 69, pp. 266-278.

Weatherson, B. [*ms*]: 'E ≠ K', available at <http://brian.weatherson.org/EK.pdf>

Weisberg, J. [2009]: 'Commutativity or Holism: A Dilemma for Conditionalizers', *British Journal for the Philosophy of Science*, 60, pp. 793-812.

Williamson, T. [2000]: *Knowledge and Its Limits*, Oxford: Oxford University Press.

---

[i] Note that this question loses much of its interest if one thinks that evidence is non-propositional. For the bulk of this paper, I make the simplifying assumption that all evidence is propositional. There is a substantive issue here. If all evidence is propositional then, strictly speaking, *experiences* can't be evidence. Williamson ([2000], pp. 194-9.) offers arguments on behalf of this assumption. Conee & Feldman ([2008]) argue for a view that has experiences as evidence; they also respond to Williamson's arguments (pp. 100-4). Though I adopt the assumption for the bulk of this paper, it is relaxed in Section 6.2.

[ii] There are Bayesian views that reject or modify Conditionalization. I focus on Conditionalization since it is the most common updating rule and because Williamson endorses a rule very similar to it. At the end of the paper I consider Richard Jeffrey's ([1965]) alternative to Conditionalization.

[iii] As I show at the end of Section 4, a slight variant of the argument goes through even against someone who denies that the Bayesian view is one that recognizes inferences at all.

[iv] There need be no fundamental distinction between these for the distinction to be drawn. For instance, one can draw this distinction even if one thinks that deductive inferences are merely limiting cases of inductive inferences: the line is then drawn between the limiting cases and all the rest.

[v] Peter Klein ([1981]) may hold a view that is amenable to this kind of response. For details, see the analysis of Klein's view in Dodd ([2007]), pp. 640-1.

[vi] Weatherson ([*ms*]) argues against Williamson's E=K thesis based on the fact that E=K allows inferential evidence. However, Weatherson does not show that inferential evidence is incompatible with the Standard Bayesian framework, rather he relies on a particular intuitive case. Williamson could respond to Weatherson's argument by simply denying Weatherson's intuition about this case. This option is not available with the argument here.

[vii] Notice that on this construal intermediate evidence and ultimate evidence are not exclusive. For my purpose, this is not important, but one could easily alter the definitions I offer in the main text to ensure that no ultimate evidence also counts as intermediate evidence.

[viii] David Christensen ([1994]) discusses the virtues of this kind of response to cases similar to—though not identical to—Bird's case.

[ix] Notice that this is a way of vindicating the classic Bayesian view discussed in Section 1, which sees evidence as corresponding to those propositions one acquires via direct perception. A response similar to the one given to Bird could be given to my example in Section 1 of the proposition that thus sun has risen being evidence.

[x] Note that in this example E is given full credence. The same story could be told if E were to receive less than full credence, but this needlessly complicates the presentation.

[xi] It can be shown that this model has interesting consequences for scenarios where agents *lose* evidence (scenarios like Bird's poisonous liquid case). There are important questions about Jeffrey Conditionalization, in particular, the looming problem of non-commutativity first discussed in Field ([1978]). On this issue, see Doring ([1999]), Lange ([2000]), Wagner ([2002], [2010]) and Weisberg ([2009]).