

Your Evidence is the Set of Facts that Are Manifest to You

Ram Neta
University of North Carolina, Chapel Hill

(URN)

An opaque urn contains a great many balls. You don't know precisely how many, but you do know that there are dozens, and that each ball is either completely black or completely white. Because the urn is opaque, you can see the balls only when they are removed from the urn. Balls are removed and then replaced, one by one, so that you never see more than a single ball at once, and so that the number of balls in the urn stays constant (plus or minus the single ball being removed at any particular moment). You see 10 such removals in a row, and remember them all clearly and accurately. On each one of those 10 occasions, you saw very clearly that the ball that was removed was black. You are now asked how likely it is that

(H) the ball removed on the 11th drawing will also be black.

How do you determine the answer to the question of how likely it is that H is true?

To answer it, you will need to do at least two things. First, you will need to figure out what evidence is in your possession. And second, you will need to figure out how likely H is to be true, given that totality of evidence. Most of the efforts of confirmation theory have been dedicated to understanding the second of these two tasks, i.e., explaining the confirmation relations between any arbitrary evidence set and any arbitrary hypothesis. But in the past two decades, many epistemologists have tried to understand the first of the two tasks, viz., determining what evidence is in your possession.

But what is it for a particular bit of evidence to be "in your possession"? What is it for evidence to be evidence that you *have*? That will be my topic in this paper. In an earlier paper¹, I answered this question by appeal to the two normative roles that an agent's evidence plays in that agent's cognitive life: there is what I called the "local role" of evidence – it can serve as the agent's reason for believing something or other – and there is also what I called the "global role" of evidence – the totality of agent's evidence set at a time determines which coherent credence functions are rationally available to that agent at that time. While I continue to accept that view as correct, the view itself is consistent with various other views about what it is for an agent to have some evidence. In this paper, I will appeal to the two normative roles of evidence to argue that every element of an agent's evidence set must be a fact, and it must be a fact that the agent is – in a sense that I will spell out more fully below – in a position to *know non-inferentially*, even though she need not already know it, or even believe it.

¹ Neta 2008.

In discussing this topic, I will return later to the example **URN** above. I will also assume what I have argued for elsewhere², viz., that the elements of an agent's evidence set are individuated as finely as propositions are individuated.

Also, I will develop my own argument by building on the insights contained on both sides of a debate between Williamson and Goldman about the nature of evidence. Specifically, I will argue that, although Goldman successfully criticizes some of the features of Williamson's account of evidence, Williamson is right that our evidence must be true. And I will also argue that, though Williamson is right about that, Goldman is nonetheless right that, for a proposition p to be in an agent's evidence set at a time t , the agent must have non-inferential proposition justification for p at t . My task will be to develop an account of evidence that is not only consistent with, but also explanatory of the truth of, both of these insights.

I. Williamson's defense of E=K

Williamson 2000 argues for an equation of an agent's evidence set with the set of all and only those propositions that the agent knows to be true. He does this by first arguing that knowledge is necessary for evidence, and then by arguing that it is sufficient for evidence.

Williamson's argument for the necessity claim proceeds by considering a case like **URN**, in which you draw 10 balls, see clearly that each ball is black, and you remember each drawing perfectly. Now, consider two hypotheses:

h: Draws 1 to 10 were black; draw 11 is white.

h*: Draw 1 was white; draws 2 to 11 are all black.

Of such a case, Williamson says, "it is natural to say that h is consistent with my evidence and that h* is not."³ In particular, Williamson elaborates, it is consistent with my evidence that draw 11 is white; it is not consistent with my evidence that draw 1 is white. Thus my evidence does not include the proposition that draw 11 is black. Why not? Given the evidence I've received from the first 10 draws, I am justified in believing that draw 11 is black, and we can stipulate that I do in fact believe this justifiably, on the basis of this evidence. Furthermore, we can stipulate that draw 11 is in fact black, so that my belief is not only justified, but also true. Why then, is the proposition that draw 11 will be black not in my evidence?

² *Ibid.* I criticize the argument of Williamson 2000 that the elements of an agent's evidence set are propositions, but I offer a different argument for the conclusion that those elements must be individuated as finely as propositions. Roughly, unless these elements are individuated at least as finely as this, it will be impossible for someone to have evidence that makes it rational for her to believe that she is seeing the evening star without also making it rational for her to believe that she is seeing the morning star: but this is obviously possible.

³ Williamson 2000, 201.

Williamson thinks that there is only one plausible answer to this question: namely, I do not know that draw 11 is black. While I might justifiably believe, on the basis of the first 10 draws being black, that draw 11 is also black, this justified belief is not knowledgeably held.

This is an argument for the conclusion that a proposition is in an agent's evidence set at a time *only if* the agent knows that proposition to be true. A proposition can be both true and justifiably believed – as is the proposition that draw 11 is black – and still fail to be in one's evidence set, so long as it fails to be known.

Now, here is Williamson's argument for the claim that knowledge is sufficient for evidence:

“Any restriction on what counts as evidence should be well-motivated by the function of evidence. [By the argument above] one's evidence includes only propositions which one knows. If, when assessing an hypothesis, one knows something *e* which bears on its truth, should not *e* be part of one's evidence? Would it not violate the total evidence condition to do otherwise?” (Williamson 2000, 203 – 4)

Williamson then criticizes two potential responses to this question: first, treating all one's knowledge as evidence would obliterate evidential interconnections among the various things one knows, and second, treating all knowledge as evidence makes inexplicable the susceptibility of our knowledge to be undermined by further evidence. Let's consider these responses one by one.

So first, if all our knowledge is our evidence, then how can some of the things we know (e.g., that 13 of the past 15 years have been the hottest years on record) serve as evidence for other things that we know (e.g., that sea levels will rise)?

Williamson answers this question by appeal to the following view concerning the “evidence for” relation:

EV *e* is evidence for *h* for *S* if and only if *S*'s evidence includes *e* and $P(h|e) > P(h)$.⁴

To see how this works, consider the case in which

e = the first 9 draws were all black,

h = the 10th draw was black,

P is the agent's epistemic probability function before acquiring knowledge that *h* is true.

⁴ Williamson 2000, 187. Williamson's acceptance of EV is in keeping with Bayesian orthodoxy. See Fitelson 1999 (especially footnote 2, which provides a long list of distinguished proponents of one or another version of EV). A compelling argument against EV is given in Achinstein 2001, and discussed below in the text.

In this case, $P(h|e) > P(e)$. In other words, the first 9 draws being black make it more likely than it was prior to those 9 draws that the 10th draw will be black. Thus, EV implies that, in this particular case, e is evidence for h . But this is consistent with e and h both being known in URN. Thus, Williamson concludes, evidential interconnections between different things that one knows are not obliterated by treating all of one's knowledge as one's evidence. Some of the things that one knows can still be evidence for other things that one knows.

Next, if all our knowledge is our evidence, then how can we explain the susceptibility of our knowledge to be undermined by further evidence? To see more clearly what the problem seems to be here, consider the rule of conditionalization, which says that one should update one's probability function by setting one's new unconditional probability in a proposition to one's prior probability of that same proposition, conditional on the evidence that one just received.

$P_{\text{new}}(h) = P_{\text{old}}(h|e)$, where the change that takes place between "new" and "old" is simply the acquisition of e into one's evidence set, and P_t is the agent's probability function given *all* of her evidence at time t .

But this rule cannot result in the diminution of $P_t(e)$ below 1 at any time $t' > t$, if $P_t(e) = 1$. In other words, conditionalization can never result in a proposition's ceasing to be evidence, once it has become evidence.

Of course, this is not to say that a proposition can never cease to be part of an agent's evidence: the agent could die, or simply forget. But it is to say that no such change could occur as a result of the application of conditionalization. Thus, if conditionalization is the only way that an agent can *rationally* revise her probability function over time, then no rational revision of an agent's probability function can result in the loss of evidence. If all knowledge is evidence, then no rational revision of an agent's probability function can result in the loss of knowledge. And Williamson regards that consequence as implausible: of course it may sometimes **be** rational for us to cease to believe a proposition that we previously knew to be true, at least if we receive a great deal of misleading evidence against that proposition. Here is Williamson's example of such a case:

"I put exactly one red ball and one black ball into an empty bag, and will make draws with replacement. Let h be the proposition that I put a black ball into the bag, and e the proposition that the first ten thousands draws are all red. I know h by a standard combination of perception and memory, because I saw that the ball was black as I put it into the bag a moment ago. Nevertheless, if after ten thousand draws I learn e , I shall have ceased to know h , because the evidence which I shall then have will make it too likely that I was somehow confused about the colours of the balls." (Williamson 2000, 205)

This example is supposed to show that what an agent knows at one point is something that the agent may be rationally required to cease believing at a later

time, simply through the acquisition of new evidence. But if conditionalization is the only way in which rationality requires us to update our epistemic probability function, then such defeasible knowledge cannot be included in an agent's evidence set. Thus, we might think, not all our knowledge can be included in our evidence.

But this argument against the sufficiency of a proposition's being known to be true for its being an agent's evidence suffers from one major defect, according to Williamson: the defect is its assumption that conditionalization is the only rational way to update our epistemic probability function. Williamson takes his example of the red ball and black ball to tell against this assumption, just as it tells in favor of the defeasibility of knowledge. As Williamson puts it: "on any reasonable theory of evidence, an empirical proposition which now counts as evidence can subsequently lose its status as evidence without any forgetting, if future evidence casts sufficient doubt on it." (Williamson 2000, 206)

Thus, neither the worry about evidential interconnections, nor the worry about conditionalization, motivates a restriction on the kinds of knowledge that can be an agent's evidence. Williamson concludes that knowledge is not only necessary for evidence, but also sufficient for it. Thus, Williamson concludes, the set of propositions that an agent knows to be true, and the set of propositions that are an agent's evidence, have the same elements, and are therefore the same set. $E = K$.

II. Goldman's defense of $E = NPJ$

Goldman 2009 criticizes Williamson's argument for $E = K$. Specifically, Goldman criticizes Williamson's argument for the necessity of knowledge for evidence, and Goldman argues that the sorts of considerations that Williamson offers for the necessity of knowledge for evidence actually support an alternative conception of evidence, according to which a proposition is in an agent's evidence set at a time if and only if that proposition is non-inferentially propositionally justified for the agent at that time. If the latter conception is correct, then, as I will show in this section, it provides Goldman with a way to rebut Williamson's argument for the sufficiency of knowledge for evidence as well.

Let's review Williamson's considerations again, one by one, to see whether those consideration support $E = NPJ$ at least as strongly as they support $E = K$.

Recall Williamson's argument for the necessity of knowledge for evidence: If I have seen that the first 10 draws of balls from an urn produced black balls, then I can reason to the likely conclusion that the next ball drawn will also be black. But while it may be true, and I may justifiably believe, that the next ball will be black, it is nonetheless consistent with my evidence that the next ball drawn is not black. Therefore, an agent's justified true belief that p is not sufficient for p 's being in that agent's evidence set. Nothing short of knowing that p will suffice for p to be in that agent's evidence set.

Goldman's reply to this argument is straightforward: all that the hypothetical case shows is that an agent's *inferentially* justified true belief that p is not sufficient for p's being in an agent's evidence set. The case shows nothing about whether an agent's *non-inferentially* justified true belief that p suffices for p to be in the agent's evidence set. Indeed, the structure of the case – in which a true proposition is justifiably believed on inductive grounds, yet not included in the believer's evidence set – seems to threaten $E = K$ as well, at least if we allow that we can sometimes arrive at knowledge on the basis of inductive reasoning from our evidence. If that is allowed, then consider a variation on URN in which the body of evidence is sufficient to provide inductive knowledge of the color of the next ball drawn: in such a case, does the agent's acquisition of this inductive knowledge on the basis of one body of evidence expand the agent's evidence to include the inductively known fact? Quite clearly not: while my inductive evidence might give me knowledge that h is true, it doesn't follow that h thereby becomes part of my evidence. Part of the point of calling my evidence for h "inductive" is that the negation of h is consistent with my total evidence, and thus h cannot itself be an element in my evidence set.

Williamson has therefore not made his case that a proposition must be known in order to be evidence: as Goldman points out, an even better explanation of the data in URN is that a proposition must *be non-inferentially* known – or at least one for which the agent has non-inferential justification to believe – in order to be evidence.

What about Williamson's argument for the sufficiency of knowledge for evidence? Can Goldman rebut those arguments as well? Although he doesn't attempt to do so, I think it's clear that he can. Recall [that](#) Williamson's argument for sufficiency comes in two steps. The first step is a rhetorical question:

"If, when assessing an hypothesis, one knows something e which bears on its truth, should not e be part of one's evidence? Would it not violate the total evidence condition to do otherwise?"

The second step is an attempt to rebut arguments (concerning evidential interconnections, on the one hand, and the indefeasibility of knowledge through conditionalization, on the other) that purport to show that the rhetorical question has a negative answer.

The problem with this argument for sufficiency is, first, that the rhetorical question doesn't even appear to do the job that Williamson wants it to do, and second, both of his rebuttals are unsuccessful. Let's go through these points one by one.

First, compare the rhetorical question that Williamson asks with others that he could have also asked, with no difference in plausibility:

"If, when assessing an hypothesis, one justifiably believes something e which bears on its truth, should not e be part of one's evidence? Would it not violate the total evidence condition to do otherwise?"

“If, when assessing an hypothesis, one is justified in believing something e which bears on its truth, should not e be part of one’s evidence? Would it not violate the total evidence condition to do otherwise?”

“If, when assessing an hypothesis, it is rationally incumbent upon one to be confident of something e which bears on its truth, should not e be part of one’s evidence? Would it not violate the total evidence condition to do otherwise?”

These different questions specify different conditions as sufficient for evidence, and yet no one of these questions is clearly more or less rhetorically effective in inviting an affirmative response than any other. This indicates that Williamson’s rhetorical question can’t do its intended argumentative work in recommending knowledge, rather than some other condition, as sufficient for evidence.

Now, what about Williamson’s argument that, even if knowledge suffices for evidence, this neither erases evidential interconnections among things that one knows, nor does it prevent knowledge from being defeasible? I’ll assess each argument in turn.

Williamson’s argument for the conclusion that $E = K$ does not erase evidential interconnections among the things that one knows depends on his undefended principle

EV e is evidence for h for S if and only if S ’s evidence includes e and $P(h|e) > P(h)$.

But EV is false. Achinstein 2001 presents several counterexamples to EV. Here is one such example. Suppose you have a headache, and you know that there is an 85% chance that your headache will go away by Noon. Still, you are impatient, so you take pill X to relieve your headache. You do this because you know that pill X has a 95% chance of relieving your headache by Noon. If this is all the relevant information you have, then you should have credence .95 that your headache will be relieved by Noon. But as soon as you form this credence, you learn that pill X is occasionally lethal, and so you become worried. You learn that pill Y will completely neutralize the effects of pill X, but you also learn that pill Y will have a 80% chance of relieving your headache by Noon. Knowing these two things, you decide to take pill Y. Now, you’re no longer worried about the potentially lethal consequences of pill X, but you should lower your credence that your headache will be relieved by Noon from .95 to .80. In this case, the fact that you took pill Y is evidence that your headache will be relieved by Noon, but it also lowers the probability that your headache will be relieved by Noon. Thus, not only does evidence for a hypothesis sometimes not raise the probability of that hypothesis – it sometimes lowers the probability of that hypothesis. EV is therefore false.

Notice also that, even if EV could be defended from such objections, the most that Williamson could use EV to establish is the compatibility of $E = K$ with evidential

interconnections. But, as Williamson notes, the evidential interconnections established by EV are symmetric. (If $P(h|e) > P(h)$, then $P(e|h) > P(e)$.)⁵ And yet it seems that the evidential interconnections that obtain among the various things we know are not all symmetric: some of those evidential connections are antisymmetric dependencies. Consider, for instance, my auditory knowledge that there is a knock on my office door, and my knowledge that a person wants to talk to me. These two pieces of knowledge are evidentially interconnected, but not symmetrically. The first is evidence for the second, but – since people are constantly emailing or calling or texting me wanting to talk to me – the second is no evidence at all for the first.

Finally, let's turn to Williamson's argument that $E = K$ is consistent with the defeasibility of knowledge by additional evidence. Williamson's case is one in which the acquisition of new evidence defeats knowledge by making it impossible to retain old evidence. For instance, after having drawn a red ball 10,000 times in a row, it is no longer rational for me to believe what I once knew to be true, viz., that I had put a black ball into this bag before making these 10,000 draws. Is this a case in which acquiring new evidence is constitutively (and not merely causally) connected to my losing old evidence?

No. To see this, consider what would happen if the bag were transparent, so I could continue to see the black ball in the bag throughout all the draws.⁶ In such a case, while I might certainly be puzzled about why the first 10,000 draws each produced a red ball, that series of draws would give me no reason to doubt that I had put a black ball in the bag: I could continue seeing it in the bag throughout all the draws, just as I had done at the beginning. In other words, the 10,000 red draws make it rational for me to doubt the proposition that was in my evidence set prior to the draws because they make it rational for me to doubt that the proposition was one that I had *observed* (and thereby been non-inferentially justified in believing) to be true. So long as I continue to observe that same proposition to be true, no number of red draws can make it rational for me to doubt that I put a black ball in the bag.

This is not to imply that observation is the only way in which I can obtain non-inferential justification for believing some proposition – perhaps, in some cases, memory can provide such non-inferential justification as well. But the only cases in which memory can provide such justification is when the accuracy of the memory itself is not called into question. Once the accuracy of a memory is called into question, that memory can no longer provide me with non-inferential justification for believing anything – if I continue to have justification for belief, such justification must now be inferential, and partly based on whatever reasons I have to trust my memory despite the factors that call its accuracy into question.

⁵ Williamson 2000, 204.

⁶ This thought experiment is from Neta 2009.

Williamson may reply that what I've said just now in effect concedes his main point about the rational defeasibility of evidence upon the acquisition of further evidence. If remembering that p can make it the case that p is in my evidence set, but then future evidence can call into question the accuracy of that memory, and thereby remove p from my evidence set, isn't that in effect to allow precisely the rational defeasibility of evidence, and so the rationally permissible violations of conditionalization, that Williamson claims?

Not exactly. Let's notice first that Bayesian conditionalization is a special case of a more general updating rule that Richard Jeffrey defended, and that is today known as "Jeffrey conditionalization":

$P_f(H) = P_i(H/E) \times P_f(E) + P_i(H/\sim E) \times P_f(\sim E)$ [where the prior probabilities of E and H are both greater than zero]

In other words, when you acquire evidence E , then set your credence in hypothesis H equal to the conditional probability of H on E , multiplied by the probability of E , added to the conditional probability of H on $\sim E$, multiplied by the probability of $\sim E$. In the special case where E is completely certain, the second addend equals 0 and Jeffrey conditionalization just becomes Bayesian conditionalization. But that is true only for the special (and rare) case where E is completely certain. Most of the propositions in our evidence set are less than completely certain.

What Jeffrey showed is that violations of the rule of Jeffrey conditionalization involve commitments to the fairness of a dynamic Dutch Book, and are therefore irrational. Jeffrey conditionalization becomes Bayesian conditionalization for the special case in which E has probability 1, i.e., in which it is rational to be certain that E is true. But how often is it rational to be certain that a particular evidence proposition is true? Williamson's own answer to this question is very plausible: not that often. While it is rational for us to be highly confident of propositions in our evidence set, it is almost never rational for us to be completely certain of those propositions.

But if Williamson admits that it is not rational for us to be completely certain of the truth of our evidence propositions, why then does he say that our evidence propositions have probability 1? For Williamson, the epistemic probability of a proposition is its probability conditional on our total evidence, and this will of course be 1 for any proposition that's an element of our total evidence. But if these probability values for propositions are not intended to be the values that our credence function is rationally constrained to assign to propositions, then on what grounds should we think that these values form a probability function at all? As Kaplan 2009 has argued, Williamson doesn't answer this question, and it's not clear what answer he could offer.

So, putting aside the question of why Williamson thinks there is any such as "epistemic probability", let's return to the main issue here: is $E = K$ consistent with

the defeasibility of knowledge by additional evidence? Williamson's argument for an affirmative answer to this question is vitiated by a poor choice of example – an example in which the agent doesn't retain the very same evidence that she had initially, and if she were to do so, that would change the case in such a way that it would no longer (by our intuitive lights) be a case of defeat. Of course, if we accept Jeffrey conditionalization as the unique constraint that rationality imposes on our credal updating, then that is – so far as we've seen at least – consistent both with $E = K$ and with the defeasibility of knowledge by additional evidence. But Williamson does not accept Jeffrey conditionalization as the unique constraint that rationality imposes on our credal updating. Thus, while it is not clear what, if any, specific challenge there is to the consistency of $E = K$ with the defeasibility of knowledge by additional evidence, it's quite obvious from our discussion of the hypothetical case above that there is no conflict at all between Goldman's $E = NPJ$ and the defeasibility of knowledge by additional evidence: as soon as one's initial, non-inferential justification for a proposition is defeated, any additional justification that one might get for that proposition doesn't return that proposition into one's evidence set, but rather justifies one in believing that proposition on the basis of other evidence.

In short, none of the considerations to which Williamson appeals in support of $E = K$ provide a reason to prefer $E = K$ over Goldman's $E = NPJ$.

So should we conclude that $E = NPJ$? Not so fast.

III. Why must evidence be factive?

If we want to know what it takes for a particular set of things to constitute an agent's evidence set at a time, we should begin by asking: what role is evidence supposed to play? Earlier in this paper, I mentioned two such roles: there is what I called the "local role" of evidence – it can serve as the agent's reason for believing particular propositions – and there is also what I called the "global role" of evidence – the totality of agent's evidence set at a time determines which coherent credence functions are rationally available to that agent at that time. Evidence can be my reason for believing something in particular, and it can also be my reason for having one rather than another coherent credence function. In playing each of those roles, my evidence imposes a *substantive*, and not merely a *structural*, rational constraint on my credal state. This contrast between substantive and structural rational constraints is sometimes thought of in terms of "narrow scope" and "wide scope" rational constraints, but the distinction I mean to pick out here is fundamentally metaphysical rather than grammatical. A particular bit of evidence imposes a substantive rational constraint on an agent's doxastic relation to particular propositions in the sense that it imposes a rational requirement on an agent to adopt a more or less specific doxastic relation to those propositions. The totality of an agent's evidence imposes a substantive rational constraint on an agent's distribution of confidence across all propositions in the sense that it imposes a rational requirement on an agent to adopt a more or less specific coherent credence function. (The qualifier "more or less specific" in both of the previous two sentence

is intended to leave room for different views on the issue of whether epistemic rationality is permissive.) In short, *for an agent to have evidence E is for E to impose substantive, rational constraints on an agent's credal state.*

But what else, besides evidence, imposes substantive, rational constraints on an agent's credal state? The demands of logical consistency and probabilistic coherence impose structural rational constraints on agent's credal state. An agent's needs or well-being may impose substantive, rational constraints on her desiderative state. And a variety of factors, including the agent's evidence as well as features of the agent's psychological processing, may impose substantive, rational constraints on an agent's activity of reasoning. But nothing else besides an agent's evidence imposes substantive, rational constraints on an agent's credal state at a time. I conclude that the following view of evidence is correct:

What is in an agent's evidence at a time is just whatever is the source of the substantive, rational constraints on the agent's credal state at that time.

This feature of an agent's evidence provides us with a way of figuring out the features of an agent's evidence, viz., by thinking about the factors that determine the substantive, rational constraints on an agent's credal state. We can apply the account I've just offered – the account of evidence as the source of substantive rational constraints on an agent's credal state – to thinking through another issue on which Williamson and Goldman disagree: the issue of whether an agent's evidence must be true. If Williamson is right that $E = K$, then the answer to this question is yes – only true propositions can be known to be true. If, however, Goldman is right that $E = NPJ$, then the answer to this question is no – even false propositions can be non-inferentially justified for an agent. In the preceding section, I defended Goldman's argument that the considerations that Williamson adduces in defense of $E = K$ tell at least as strongly in favor of $E = NPJ$. Does this indicate that, as Goldman would have it, evidence need not be factive? In the remainder of this section, I will argue for a negative answer to this question. More specifically, I will argue that, since evidence is the source of substantive rational constraints on an agent's credal state, an agent's evidence must consist solely of *facts*, or truths. An agent's evidence cannot, I will now argue, be false.

Let me begin by considering Williamson's own defense of the view that an agent's evidence must be true. His defense consists in two parts. First:

“That propositional evidence is knowledge entails that propositional evidence is true. That is intuitively plausible; if one's evidence included falsehoods, it would rule out some truths, by being inconsistent with them. One's evidence may make some truths improbable, but it should not exclude any outright. Although we may treat false propositions as evidence, it does not follow that they are evidence. No true proposition is inconsistent with my evidence, although I may think that it is.”⁷

⁷ Williamson 2000, 201.

In other words: while it's possible for falsehoods to be probable on my evidence, it's impossible for a falsehood to be entailed by my evidence. Thus, my evidence must be true.

And second:

“adjusting our beliefs to the evidence... is a way of adjusting them to the truth. Although true evidence can still support false conclusions, it will tend to support truths.”⁸

In other words: the best explanation of why I should adjust my beliefs to the evidence is that, by so doing, I am adjusting my beliefs to the truth.

Goldman raises important challenges to each of these arguments. In response to the first argument, he wonders why evidence cannot be deductively as well as inductively misleading. In response to the second argument, he wonders why the best explanation of why I should adjust my beliefs to the evidence isn't rather that, by so doing, I am adjusting my beliefs to propositions that are likely to be true, and so adjusting my beliefs in a way that is likely to lead to the truth.⁹

In the remainder of this section, I will argue that Goldman's challenges fail: if Goldman were right that $E = NPJ$, then we would be left with no way of understanding the notion of “likelihood” in such a way that we could explain why adjusting my beliefs to the evidence is adjusting them in a way that it likely to lead to the truth. Specifically, I will argue that, while there are ways of understanding “likelihood” so that adjusting my beliefs to the evidence is adjusting them to propositions that are likely to be true, and there are also ways of understanding “likelihood” so that adjusting my beliefs to propositions that are likely to be true is adjusting my beliefs in ways that are likely to lead to the truth, there is no interpretation of “likelihood” on which both of these two claims hold, and so no interpretation of “likelihood” on which Goldman can explain the point of adjusting my beliefs to the evidence. Since Goldman offers no explanation of this fact, and it's not clear how a defender of $E = NPJ$ could explain this fact, we should reject $E = NPJ$ in favor of a view that requires evidence to be true. That is the outline of my argument in the remainder of this section. The next section describes a view of evidence similar to Goldman's $E = NPJ$, but on which evidence requires truth as well. For now, we can think of it as the view that your evidence is all and only those *facts* that are non-inferentially propositionally justified for you. In the next section, I will try to make clear why this view is not as gerrymandered as it may sound, and explain how this view explains why evidence is the source of all substantive rational constraints on our credal states.

⁸ Williamson 2000, 202.

⁹ Goldman 2009, 88 – 9.

Our question now is: if $E = NPJ$, then why think that, when I adjust my beliefs to my evidence, I am adjusting them in a way that is likely to lead to the truth? Goldman answers this question by saying that my evidence is *likely* to be true, and so adjusting my beliefs to my evidence is adjusting them to propositions that are likely to be true. But what does this claim amount to? Recall, on Goldman's view, my evidence propositions are all propositions that I am justified in believing – and for me to be justified in believing those propositions is for those propositions to be likely (from my point of view, suitably understood) to be true.

The parenthetical qualifier in the preceding formulation is needed in order to identify the kind of likelihood of truth that's involved in a proposition's being justified for me. A proposition could be likely to be true because, say, the proposition concerns some future state of the universe, and the current state of the universe makes it objectively highly likely that this future state will obtain. But such objective likelihood has nothing to do with whether or not a proposition is justified *for a particular agent*. A proposition is justified for a particular agent only if it is, *from that particular agent's point of view*, likely to be true – whether or not the proposition is objectively likely to be true, given some prior state of the universe. But how are we to understand the distinction between objectively likely truth, on the one hand, and likelihood of truth *from some particular point of view*, on the other?

We may think to draw this distinction is by appeal to the agent's beliefs: a proposition P is likely to be true from a particular agent's point of view just in case P is objectively likely to be true conditional on the truth of the agent's beliefs. Unfortunately, this way of drawing the distinction cannot succeed, since most agents' beliefs are inconsistent, and so those beliefs will entail everything, and therefore the conditional probability of *any* proposition on the truth of the agent's beliefs will be the same, namely, 1.

Another way we might try to distinguish objectively likely truth from likelihood of truth from some particular point of view is by appeal to the agent's knowledge: a proposition P is likely to be true from a particular agent's point of view just in case P is objectively likely to be true conditional on the agent's knowledge. Unfortunately, this way of drawing the distinction cannot serve Goldman's purposes, for many of the propositions that we are non-inferentially propositionally justified in believing do not have a high probability conditional on what we know. Suppose, for instance, I am looking at playing cards being displayed for a half-second each, one at a time: a 4 of Hearts, a King of Spades, an Ace of Clubs, a 7 of black Diamonds – but, despite seeing the 7 of black Diamonds, I fail to recognize it as a 7 of black Diamonds, and find myself confused about what it is. Although I am visually presented with a 7 of black Diamonds, and so am non-inferentially propositionally justified in believing that it is a 7 of black Diamonds, I do not know any facts conditional on which it is likely that what I see is a 7 of black Diamonds. Thus, what I am non-inferentially (propositionally) justified in believing does not always have a high probability

conditional on what I know. And so Goldman cannot explain the sense in which my evidence is likely to be true, on the supposition that $E = NPJ$.

Yet another way we might try to distinguish objectively likely truth from likelihood of truth from some particular point of view is by appeal to the fact that the agent has specific beliefs: a proposition P is likely to be true from a particular agent's point of view just in case P is objectively likely to be true conditional on the agent's having precisely the beliefs that she happens to have (i.e., the proposition on which the probability is conditional is a proposition specifying the content of the agent's actual beliefs). Unfortunately, this way of drawing the distinction cannot succeed, since many of the propositions that are non-inferentially justified for an agent are not likely to be true conditional on the agent's having the beliefs that she has. Again, consider the case of the 7 of black Diamonds, described above: my evidence includes the proposition that there is a 7 of black Diamonds, but that proposition is not likely to be true given my having any of the beliefs that I have.

Finally, another way we might try to distinguish objectively likely truth from likelihood of truth from some particular point of view is by appeal to the agent's evidence: a proposition P is likely to be true from a particular agent's point of view just in case P is objectively likely to be true conditional on the agent's evidence. Of course, if this is how Goldman understands the notion of likelihood of truth from the agent's point of view, then it follows trivially that every proposition in an agent's evidence set is likely to be true – in fact, each such proposition is maximally likely to be true. Could this account of the likelihood of our evidence from our own point of view provide Goldman with the explanation that he needs? No, and here's why. Although this account of likelihood from a particular point of view entails Goldman's view that an agent's evidence is likely, from her point of view, to be true, it leaves it completely unclear why we should now accept the claim that adjusting my beliefs to my evidence – i.e., to propositions that are likely to be true on my evidence – is adjusting my beliefs in a way that is likely to lead to the truth. The view that my evidence propositions are likely (from my point of view) to be true because their conditional probability on themselves is high does nothing to explain why adjusting my beliefs to my evidence propositions is more likely to lead to the truth than, say, adjusting my beliefs to the negations of my evidence propositions.

Of course, Goldman could insist that the reason I'm more likely to get to the truth by adjusting my beliefs to my evidence propositions rather than to their negations is because my evidence propositions are justified for me but their negations are not. But why does this difference in justification explain the difference in likelihood of truth to which Goldman appeals? Our goal was to try to understand the sense in which, on Goldman's view, adjusting my beliefs to the evidence was adjusting my beliefs to propositions that are likely to be true, and so adjusting my beliefs in a way that is likely to lead to the truth. This appeal to justification does not explain the sense in which it is both likely that my evidence propositions are true and also likely that adjusting my beliefs to my evidence is adjusting them in a way that is likely to lead to the truth. And unless he can identify a particular notion of likelihood that is

able to do both of these jobs, Goldman's explanation of the point of adjusting my beliefs to the evidence fails.

Since we have not been able to identify such a notion of likelihood on Goldman's behalf, I conclude that his explanation fails, and that Williamson's account of evidence does enjoy one significant advantage over Goldman's: Williamson, unlike Goldman, can explain the point of adjusting our beliefs to our evidence.

Is there an account of evidence that can do all the work that Goldman's account can do, while also explaining the point of adjusting our beliefs to our evidence? In the next section, I specify such an account.

IV. Your Evidence is What Is Manifest to You

Goldman is right that, for a proposition p to be in an agent's evidence set at a time t , the agent must have non-inferential proposition justification for p at t . But Williamson is right that our evidence must be true. Is there an account of evidence that explains both of these necessary conditions, without also entailing further conditions that are not in fact necessary?

We might gerrymander such an account: your evidence is all and only those truths that are non-inferentially propositionally justified for you. But, while this account might get the extension right, the fact that it is gerrymandered suggests that it cannot explain why it gets the extension right. An adequate account of an agent's evidence should be not just extensionally adequate, but also explanatory. In this respect, the present account of evidence is no better than the stipulative account that David Lewis 1996 offers of your evidence as simply your perceptual experiences and memories: the account may be extensionally adequate, but it does not explain why.

Is there an account that can explain why your evidence includes all and only those facts that are non-inferentially propositionally justified for you? I believe there is, and that the account makes appeal to the notion of a fact being, as McDowell 1995 says, "manifest". A fact p is manifest to an agent A if the very fact that p constitutes A 's propositional justification for p . Let me give some examples:

What is it that currently constitutes your propositional justification for "I exist", "I am conscious", "I am thinking about epistemology"? It seems clear that, in each case, what constitutes your justification for each of these facts is simply the fact itself. Of course, the fact that you are conscious will constitute *your* justification for that proposition without constituting *my* justification for that proposition – since my justification for that proposition, if any, will rely on whatever evidence I may have about your consciousness. Since the facts that I exist, that I am conscious, and that I am thinking about epistemology, all constitute my propositional justification for the

corresponding propositions, it follows, on the present account of evidence, that all are in my evidence set.

My evidence set will therefore include all and only those facts that are justified for me in the same (non-inferential) way that these facts are justified for me. Will all such facts be facts about my own consciousness? If the answer to this question were yes, then the phenomenalist conceptions of evidence that were popular nearly a century ago would have been correct. But, as I have argued elsewhere¹⁰, the answer to this question is no, despite the contextual pressures¹¹ that sometimes lead philosophers to think that the answer is yes.

The facts that are non-inferentially justified for me are simply the facts that I am in a position to know non-inferentially. I might fail to know them because I fail to believe what I am justified in believing. Or I might fail to know them because my beliefs are not properly based. But, if a fact is one that I am justified in believing, then, whether or not I believe it, there is at least one clear sense in which I am in a position to know it. Because my evidence set includes all and only those facts that are non-inferentially justified for me, and the facts that are non-inferentially justified for me are simply the facts that I am in a position to know non-inferentially, it follows that, on my account, my evidence includes all and only those facts that I am in a position to know non-inferentially.

Let me conclude by considering how this account of evidence helps to explain why evidence is the source of all the substantive, rational constraints on my credal state. Because my evidence includes those facts that are non-inferentially justified for me, they serve as *sources*, rather than transmitters, of rational constraint: if a fact is in my evidence set, then at least one thing that justifies me in believing that fact to obtain is simply: the fact itself. Of course, there are some facts that I am justified in believing inferentially, on the basis of other facts – but the former are transmitters of rational constraint, rather than sources of it.

In addition to being sources of rational constraint, though, the facts that are in my evidence set impose substantive rather than structural constraints on my credal state: when I have a certain body of evidence, I am not confronted with a choice between rationally drawing a conclusion from the evidence or rationally rejecting the evidence. My evidence is just those facts that I cannot rationally reject. The rational requirement to accept my evidence operates in tandem with the rational requirements of consistency to determine a rational requirement that I accept conclusions on the basis of my evidence. So long as such conclusions are justified inferentially though, they are transmitters of rational constraint rather than sources of it, and so they do not count as part of my evidence. This is why my evidence in

¹⁰ Neta 2007.

¹¹ Neta 2003.

URN does not include the proposition H, that the 11th draw will be a black ball, even if H is a proposition that I justifiably believe on the basis of my evidence.¹²

Works Cited

Achinstein, Peter. 2001. *The Book of Evidence*. Oxford: Oxford University Press.

Fitelson, Branden. 1999. "The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity." *Proceedings of the Philosophy of Science Association* **66**: S362 – 78.

Goldman, Alvin. 2009. "Williamson on Knowledge and Evidence" in Greenough and Pritchard, eds. 2009, 73 – 91.

Greenough, Patrick and Pritchard, Duncan, eds. 2009. *Williamson on Knowledge*. Oxford: Oxford University Press.

Hatzimoysis, Anthony, ed. 2011. *Self-Knowledge*. Oxford: Oxford University Press.

Kaplan, Mark. 2009. "Williamson's Causal Approach to Probabilism" in Greenough and Pritchard eds. 2009, 122 – 39.

Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* **74**: 549 – 67.

McDowell, John. 1995. "Knowledge and the Internal." *Philosophy and Phenomenological Research* **55**: 877 – 93.

Neta, Ram. 2003. "Contextualism and the Problem of the External World." *Philosophy and Phenomenological Research* **66**: 1 – 31.

---. 2008. "What Evidence Do You Have?" *British Journal for the Philosophy of Science* **59**: 89 – 119.

---. 2009. "Defeating the Dogma of Defeasibility" in Greenough and Pritchard, eds. 2009, 161 – 82.

---. 2011. "The Nature and Reach of Privileged Access" in Hatzimoysis ed. 2011, 9 – 32.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

¹² Thanks to Velislava Mitova for very helpful comments on an earlier draft.