

What Evidence Do You Have?

Ram Neta

ABSTRACT

Your evidence constrains your rational degrees of confidence both locally and globally. On the one hand, particular bits of evidence can boost or diminish your rational degree of confidence in various hypotheses, relative to your background information. On the other hand, epistemic rationality requires that, for any hypothesis *h*, your confidence in *h* is proportional to the support that *h* receives from your *total* evidence. Why is it that your evidence has these two epistemic powers? I argue that various proposed accounts of what it is for something to be an element of your evidence set cannot answer this question. I then propose an alternative account of what it is for something to be an element of your evidence set.

- 1 *Introduction*
 - 2 *The elements of one's evidence set are propositions*
 - 3 *Which kinds of propositions are in one's evidence set?*
 - 3.1 *Doxastic accounts of evidence*
 - 3.2 *Non-doxastic accounts of evidence*
 - 4 *Elaborating and defending the LIE*
-

1 Introduction

Your evidence constrains your epistemically rational degrees of confidence in two ways.

First, and most obviously, your evidence *locally* constrains your epistemically rational degrees of confidence, as follows: for any bit of evidence that you have, it is possible for you to use that evidence in the rational regulation of your degrees of confidence, or of your attitudes more generally. Relative to your background information, particular bits of your evidence thereby boost your rational degrees of confidence in some hypotheses, and diminish your rational degrees of confidence in others.

Second, and less obviously, your evidence *globally* constrains your epistemically rational degrees of confidence, as follows: for any hypothesis *h*, your confidence in *h* at a particular time *t* should be—as a matter of epistemic rationality—proportional to the degree to which the total evidence that you possess at *t* supports *h*. Epistemic rationality does not require that your

confidence in h be proportional to the degree to which things *other than your evidence* support h ; in fact, if the degree to which things other than your evidence support h differs from the degree to which your total evidence supports h , then epistemic rationality does not even permit that your confidence in h be proportional to the degree to which things other than your evidence support h . And epistemic rationality does not require that your confidence in h be proportional to the degree to which *only a proper part of your evidence* supports h ; again, if the degree to which that proper part of your evidence supports h differs from the degree to which your total evidence supports h , then epistemic rationality does not even permit that your confidence in h be proportional to the degree to which only a proper part of your evidence supports h . Rather, what epistemic rationality requires is that your confidence in h at t be proportional to the degree to which h is supported by *all and only* the evidence that you possess at t .

It does not follow from this requirement that one's total evidence is the only thing that gives one *justification* for believing something, or is the only thing that serves to *justify* a belief that one has, or is the only thing that determines the extent or the nature of one's justification. Indeed, from the requirement stated above about epistemic rationality, evidence, and degree of confidence, nothing whatsoever follows about epistemic justification, or about any other property of belief *simpliciter* or acceptance *simpliciter*.

Moreover, it does not follow from the requirement stated above that one is *culpable* for failing to proportion one's degree of confidence in h to the degree of support that h receives from one's total evidence. If we fail to comply with the requirements of rationality, then our state of confidence is less than fully rational. But it does not follow that we are culpable for this failure. Less than perfect rationality is one thing, blameworthiness is another, and the former may not suffice for the latter.

Furthermore, it does not follow from the requirement stated above that one's degrees of confidence should be precise, or real-valued, or conform to the Kolmogorov axioms of probability. In this paper, I will make no assumptions concerning any of these various issues.

The requirement stated above is one version of the so-called 'requirement of total evidence', and I will henceforth use that familiar phrase to refer to that requirement. I will also, in what follows, assume that this requirement of total evidence (as I'm calling it) is true. In fact, not only do I take it to be true, but I take it to be a datum with which any acceptable philosophical account of evidence must be consistent. I will, however, in the final section of this paper, respond to one important objection to this requirement.

Some philosophers may object that what I have called the requirement of total evidence implies that epistemic rationality issues *requirements* concerning how we should invest our confidence, and that this implied claim is false: in fact,

epistemic rationality merely issues *permissions* concerning how we may invest our confidence across hypotheses. I reply to this objection as follows: if there is a distribution *D* of confidence that epistemic rationality does not permit, then *D* is a distribution of confidence that epistemic rationality requires us not to have. But in that case epistemic rationality requires us to have a distribution of confidence that differs from *D*, though it might not require anything more specific than that—the requirements of epistemic rationality may, in that case, be satisfied by many different distributions of confidence (so long as each of them differs from *D*). Now, the requirement of total evidence tells us only that epistemic rationality requires that our confidence in a hypothesis be proportional to the support that our total evidence gives that hypotheses; but if the extent to which our total evidence supports that hypothesis is highly imprecise or unspecific, then the requirement of total evidence leaves open the possibility that the requirements that epistemic rationality imposes upon our confidence in that hypothesis will also be highly imprecise or unspecific. So the requirement of total evidence can accommodate what is sometimes thought of as the permissiveness of epistemic rationality. All that the requirement of total evidence says is that one's confidence in a hypothesis must be proportional to the support that that hypothesis receives from one's evidence, however imprecise or unspecific that support may be.

Some philosophers will be inclined to think that ordinary examples tell against the requirement of total evidence. Consider, say, a murder investigation. Evidence is gathered: a particular knife bears traces of the victim's blood, its handle has the culprit's fingerprints, the victim's autopsy revealed death by stabbing, the culprit had a clear motive, and so on. Should our confidence in a hypothesis concerning the identity of the perpetrator be proportioned simply to the support that this hypothesis receives from such evidence alone? It may seem not, for our rational degree of confidence in this hypothesis also depends upon our *background information*, e.g., the victim was once a living human being, knives do not simply spring into existence with peoples' fingerprints on them, and so on. Thus, the present objection concludes, we should distribute our confidence across hypotheses not in proportion to the support that those hypotheses receive from our total evidence alone, but rather in proportion to the support that those hypotheses receive from our total evidence *in conjunction with our background information*.

The problem with this objection, as it stands, is that it says nothing about *what background information is yours*.¹ And on this issue, the proponent of the requirement of total evidence can offer a suggestion: background information is yours (i.e., is admissible as helping particular bits of your evidence to determine

¹ Specifying what counts as admissible background information is not easy. For a vivid account of some of the difficulties, see Christensen [1997].

your rational degrees of confidence in hypotheses) only if that information is itself sufficiently well supported by your *total* evidence. Thus, when we conduct the murder investigation, we can assess various hypotheses in light of the evidence, against the background of such information as that the victim was once a living human being, knives do not simply spring into existence with peoples' fingerprints on them, etc. But such background information helps our evidence to determine our rational degrees of confidence in hypotheses only because that information is sufficiently supported by our total evidence (including evidence that we did not collect in the course of this particular murder investigation). If our total evidence does not sufficiently support the proposition that the victim was once a living human being, then we would not be able to use evidence to confirm or disconfirm hypotheses against a background that includes that proposition.

Now, many philosophers have attempted to understand what it is for your evidence to *support* a hypothesis *H*, or to support it *to a certain degree*. But less effort has been devoted to answering the question what it is for something to be *your* evidence, to be evidence that you *have*. Since the answer to this latter question will help to determine how you should distribute your confidence over hypotheses at a given time, it is a substantive question—a question the answer to which has substantive normative implications—what evidence you possess at a given time. We cannot answer this substantive, normatively significant question by stipulation or definition. For instance, we cannot answer it by merely stipulating that S has evidence *e* at time *t* iff *e* is (say) one of S's perceptual experiences at time *t*. No act of stipulation can answer our substantive question concerning what it is that determines our rational distribution of confidence across hypotheses at a time. Such a stipulation would simply put some pressure on us to ask our normative question using other terms besides the term 'evidence.'

So, in general, what evidence does one possess at a given time? Different philosophers have propounded different answers to this question. For instance: According to Lewis ([1996], p. 424), your evidence set contains all of your 'perceptual experiences and apparent memories.' (Your evidence set might also, Lewis allows, include the evidence of our extrasensory faculties, if they exist.)

According to Davidson ([2001], p. 141), 'nothing can count as a reason for holding a belief except another belief.' But since, as Davidson allows, any bit of evidence for *p* is a reason for holding the belief that *p*, it follows that nothing can count as a bit of evidence for *p* except a belief, presumably, a belief that one has. One's evidence set at a particular time *t* consists only of one's beliefs at *t* (whether or not it consists of all such beliefs).

According to Feldman ([1988], p. 96), your evidence set contains all and only what you are thinking about at a given time.

According to Williamson ([2000]), p is in S 's total evidence if and only if p is a proposition that S knows to be true. In other words, S knows that p if and only if p is an element of the set containing all and only the evidence that S possesses.

Which, if any, of these philosophers is right? In order to address this question, let's first consider just what it is that we demand from an answer to the question what evidence someone has. I propose the following two conditions of adequacy on any such answer.

1. *It should help to explain the piecemeal epistemic effect of each bit of our evidence.* More specifically, it should explain how a subject is able to use particular bits of evidence in the rational regulation of her own attitudes, e.g., via inference to the best explanation, probabilistic confirmation, ruling out hypotheses, or in any of the other ways in which subjects use evidence in the rational regulation of their own attitudes. (To say that subjects engage in the *rational regulation* of their attitudes is not to imply that the rational regulation of one's own attitudes is an intentional or voluntary action. There are, of course, intentional or voluntary actions that one can perform in order to affect one's rational regulation of one's own attitudes: but rational regulation is not itself—at least not typically—such an action. It is, nonetheless, an activity for which we can be held responsible, and with respect to which we can be assessed. If I fail to regulate my own attitudes in such a way as to satisfy the requirements of epistemic rationality, then I am subject to epistemic criticism, even if my failure, and the intellectual conduct that exhibits this failure, is not voluntary or intentional.)
2. *It should help to explain the global epistemic effect of the sum total of our evidence.* More specifically, it should explain why the requirement of total evidence is true. Why is it that, for any hypothesis h , your confidence in h should be proportional to the degree to which your total evidence supports h ? We can, I am supposing, agree that this is true. But we do not yet have any satisfactory explanation of *why* it is true. Why does my failure to distribute confidence in this way leave me open to rational criticism?

Williamson employs the first of these two stated conditions of adequacy to argue that the elements of one's evidence set are propositions; indeed, that they are all and only the propositions that one knows to be true. While I accept the conclusion of Williamson's argument that the elements of one's evidence set are propositions, and while I agree with Williamson that condition (1) above can be used to establish this conclusion, I do not accept Williamson's own argument from (1) to the conclusion that all evidence is propositional. Nor do

I accept Williamson's stronger conclusion that the elements of one's evidence set are all and only the propositions that one knows to be true. As we will see in Section 2 below, condition (2) can be used to argue against that stronger claim of Williamson's. So which propositions are the elements of one's evidence set? I will devote Section 3 below to answering this question.

On the basis of the conditions stated above, I will argue for the following conclusions. S's evidence set at a given time does not consist of all and only what S believes, or of what S justifiably (or rationally or blamelessly) believes, or of what S truly believes, or of what S justifiably (or rationally or blamelessly) and truly believes, or of what S knows. It also does not consist of all and only some non-propositional elements, e.g., sensory events. Rather, the proposition p is an element of S's evidence set if and only if p has the following property: if E_{conj} is a proposition which is such that, for any hypothesis h , S is rationally required to have a degree of confidence in h that is proportional to the degree to which E_{conj} supports h , then p is a member of some decomposition of E_{conj} into conjuncts each of which is such that S is able to use that conjunct in the regulation of her attitudes. In other words, the two properties of evidence that are mentioned in our two conditions of adequacy are not only *essential* properties of evidence, but they are *fundamental* as well: they jointly characterize what it is fundamentally for something to be in one's evidence set. This last, positive claim about the fundamental nature of the relation of *having evidence* requires elaboration and defense: I will devote myself to these tasks in Section 4 below.

Before proceeding, I should issue three disclaimers: First, I will say absolutely nothing in this paper about what it is for some evidence to support a hypothesis, or to support it to a certain degree. I will simply take those crucial notions for granted.

Second, some philosophers will complain that the two conditions of adequacy stated above are too weak, and an account of our evidence should do much more than simply satisfy those conditions. For instance, philosophers will want an explanation of how evidence sometimes gives us *justification* or *knowledge*, or how evidence determines how it is rational to *act*, or why it is that beliefs are more epistemically likely to be true the more strongly they are supported by our evidence, etc. And we also want an account of evidence to help us to understand our various practices of collecting, requesting, and citing evidence in the law, in science, in the humanities, and in daily life. (Our actual practices are, of course, the result of various forces all operating in tandem, including the various exigencies of daily life, and various institutional pressures upon collective practices. Of course not all of the features of these practices can be explained exclusively by appeal to the nature of evidence. But an account of evidence should at least enter into the explanation of various features of those practices.) Now, let me make this perfectly clear: *I think that such a complaint is*

entirely correct. It is true that a full account of evidence should do much more than simply satisfy the two conditions of adequacy listed above. But, as I'll argue below, just satisfying those two conditions is not easy; in fact, no account of evidence that has been proposed to date clearly does so, and neither does any easily envisaged variation upon those extant accounts. Merely to articulate an account of evidence—or even say anything informative about evidence—that does clearly and demonstrably satisfy both of those two conditions constitutes progress, however modest. That is the progress that I hope to make in this paper.

My third disclaimer is this: some philosophers will be very skeptical about the prospects of giving a fully general account of evidence, i.e., of what general kinds of things get to be the elements of our evidence set across disciplines, across contexts, across historical eras, and so on. To a large degree, I am sympathetic with this skeptical attitude. There are big and important differences between what counts for evidence in the course of an archaeological dig, in the course of a criminal trial, and in the course of an experiment in post-behaviorist developmental psychology. I do not wish to understate these differences, or to understate their importance. In this paper, I mean simply to locate some genus of which these various things are all species. We should not expect the genus itself to be anything very specific, since what all the various kinds of evidence have in common is not anything very specific. What we should expect is that the genus will be highly . . . generic.

2 The elements of one's evidence set are propositions

Williamson offers three arguments from condition (1) to the conclusion that all evidence is propositional. What he means—and what I will mean—by a 'proposition' is just the content of some belief, some knowledge, or some confidence. (I accept without question the popular view that these states all bear the same kind of content.) This characterization of propositions is intended to be neutral on various controversial issues concerning the nature of propositions: whether they are sets of possible worlds, or syntactically structured entities of some sort, or psychologically real, etc. Whatever they are, they are things that we sometimes believe to be true and sometimes believe to be false, things that we sometimes know to be true and sometimes know to be false, and things in the truth of which we are more or less confident. Sometimes what we are confident of is precisely what rationality requires that we be confident of, and sometimes it is precisely what rationality requires that we not be confident of. So propositions—whatever they are—are what rationality requires that we be more or less confident of. In this section, I will critically examine Williamson's

arguments for the propositionality of our evidence, and then offer another argument that seems to me more compelling, though still far from conclusive.

According to Williamson's first argument, subjects use evidence to choose between hypotheses by asking which hypothesis best explains the evidence. But explananda are propositional, and so evidence is propositional:

We often choose between hypotheses by asking which of them best explains our evidence[. . .]evidence is the kind of thing which hypotheses explain. But the kind of thing which hypotheses explain is propositional. [. . .] Inference to the best explanation concerns why-explanations, which can be put in the form '___ because . . .', which is ungrammatical unless declarative sentences, complements for 'that', fill both blanks. We cannot simply *explain Albania*, for 'Albania because . . .' is ill-formed. We can sometimes make sense of the injunction 'Explain Albania!', but only when the context allows us to interpret it as an injunction to explain why Albania exists, or has some distinctive feature. What follows 'why' is a declarative sentence, expressing the proposition to be explained—that Albania exists, or that it has the distinctive feature. (Williamson [2000], pp. 194–5)

In response to this first argument, one might object that, even if the conjunction 'because' can grammatically conjoin nothing other than declarative sentences, nothing about the relation of why-explanations follows from this feature of the conjunction. Might this not be a case in which grammar is metaphysically misleading?

Of course, to make this objection credible, the objector would have to provide some reason for thinking that this is a case in which grammar is metaphysically misleading; what reason could there be? Perhaps simply the obvious fact, noted by Williamson, that we commonly speak of objects as evidence. The bloody knife is evidence that the defendant is guilty. The clouds are evidence that it will rain. My sore throat is evidence that I'm coming down with a cold. Worse still, we ordinarily speak of someone as 'planting evidence,' of 'having the evidence right here on the table,' but of course one cannot plant a proposition or have a proposition on the table. These ordinary locutions suggest that the grammar of 'because,' noted by Williamson, is metaphysically misleading. Of course, it could still be that, when we speak of the bloody knife as *being evidence* that the defendant is guilty, what that amounts to is that there is some proposition that somehow involves reference to the bloody knife, and that is itself evidence that the defendant is guilty. And it could be that, when we speak of planting evidence, or of having the evidence right here on the table, what that amounts to is that there are some propositions that involve reference to objects that are planted, or are on the table, and that are themselves evidence for some contextually salient hypothesis. But if Williamson is willing to defy grammatical appearances in our account of what it is for the bloody knife to be evidence that the defendant is guilty, then why should we not be equally willing

to defy grammatical appearances when it comes to why-explanations? The considerations adduced up to now seem to leave it an open question whether the explanantia of our hypotheses are propositional, and so whether evidence is propositional.

Williamson has a second argument from condition (1) to the conclusion that evidence is propositional. This second argument depends upon the use of evidence in probabilistic reasoning. Williamson writes:

The best way of comparing the conditional probabilities of two hypotheses h and h^* on evidence e , $P(h/e)$ and $P(h^*/e)$, is often by calculating the inverse probabilities of e on h and h^* , $P(e/h)$, and $P(e/h^*)$. [...] But what has a probability is a proposition; the probability is the probability *that*. . . . At least, that is so when ‘probability’ has to do with the evidential status of beliefs, as now; if we speak in this connection of the probability of an event, we mean the probability *that it occurred*. (Williamson [2000], pp. 195–6)

Why should we accept the claim that, ‘when “probability” has to do with the evidential status of beliefs,’ then ‘what has a probability is a proposition’? Why not say instead that what has a probability is, at least in some cases, an event or a state rather than a proposition? What is the probability of the knife’s being (in the state of being) bloody, given that the defendant is guilty? What is the probability of (the event of) the sky’s clouding up, given that it will rain? What is the probability of my (being in the state of) having a sore throat, given that I am coming down with a cold? We could, of course, ask instead about the probability *that* the knife is bloody, or *that* the sky will cloud up, or *that* I have a sore throat. But why are these latter questions more important for probabilistic reasoning than the questions concerning states or events? I do not see that Williamson has given us answer to this question. But we’ll return to this question below.

Finally, Williamson has a third argument from condition (1) to the conclusion that evidence is propositional. This third argument depends upon the use of evidence to rule out hypotheses. Williamson writes:

our evidence sometimes rules out hypotheses by being inconsistent with them. . . . But only propositions can be inconsistent in the relevant sense. (Williamson [2000, p. 196])

Let’s grant that only propositions can be inconsistent in the relevant sense. Why should we allow, though, that there is an inconsistency between hypothesis and evidence itself, rather than an inconsistency between hypothesis and one or another statement of the evidence? Why not say, instead, that the hypothesis that I feel just fine is ruled out by my sore throat, and so the hypothesis is

inconsistent with a particular statement about the evidence, namely, that my throat is sore? Again, why not say that the hypothesis that the sky is bright blue is ruled out by the clouds, and so the hypothesis is inconsistent with a particular statement about the evidence, namely, that the sky is clouding up? Once again, it seems that Williamson's argument does not establish its intended conclusion, that evidence is propositional.

Is there any good reason, then, to believe that the elements of our evidence set are all propositions? Yes. In order to see what reason there is to believe this, let's begin by considering an alternative proposal:

S's evidence set at time t is the set of all and only those objects with which S is acquainted at t .

Never mind precisely what's involved in the relation of acquaintance, for there is a problem with this proposal no matter how we spell out that relation. To see what the problem is, consider the following scenario. A particular knife is bloody. The blood on the knife is—unlike most human blood—of the same unusual chemical composition as the blood of the murder victim. The knife was lying in the kitchen, which is right next to the living room, and that is where the murder victim was found. The knife had the defendant's fingerprints, and nobody else's fingerprints, all over the handle. Furthermore, there is no evidence pointing to the guilt of anyone other than the defendant.

Now, if I have enough evidence to know all of the facts just listed, then I have enough evidence to make it rational for me to be very confident that the defendant is guilty. Perhaps I don't have enough evidence to issue a verdict of 'guilty' at the trial, but at least I have enough evidence to *be very confident* that the defendant is guilty. So somehow or other, if my total evidence enables me to know all of the aforementioned facts, then that same evidence must make it rational for me to be very confident that the defendant is guilty. Let's ask then: what could my total evidence consist of, such that it makes it rational for me to be very confident that the defendant is guilty? What evidence set could render this belief rational?

Suppose my evidence is simply the set of the following objects with which I am acquainted: the knife, the corpse of the murder victim, the fingerprints, the blood, the room in which the murder victim was found, the room in which the knife was found, and so on. The objects just listed do not, by themselves, determine any answer to the question of how confident I should be about the defendant's guilt. It's not because there are too few objects; adding more objects to our list won't help. The problem is this: I can be acquainted with precisely the same objects, even if what it is rational for me to be confident of is completely different. For instance, let's consider simply the objects just listed: the knife, the blood, the kitchen, the living room, the victim's corpse, the

defendant's fingerprints. Consider another scenario in which I am acquainted with all these same objects, and come to know the following: the defendant's fingerprints are all over the kitchen, but the victim's corpse was found in the living room with blood all over it, and the knife was found free of both blood and fingerprints. If my evidence set is just that set of objects with which I am acquainted, then, since each of the two situations that I've described involves my being acquainted with precisely the same set of objects, it will be equally rational for me to be very confident that the defendant is guilty in both of the two situations. But in fact it will not be equally rational for me to be very confident that the defendant is guilty in both of the two situations. Even though the two situations involve my being acquainted with precisely the same objects, other features of the two situations give rise to my having different bodies of knowledge in the two situations, and so have different consequences for what confidence I can rationally have in the two situations. But what confidence I can rationally have is determined by my evidence set. So at least some of these other features that give rise to my having different bodies of knowledge in the two situations, also have consequences for what's in my evidence set. So my evidence set cannot be simply a set of objects with which I am acquainted. Obviously, this problem will not be solved by simply adding more objects to my evidence set. So while objects may still be elements of my evidence set, they cannot be the only elements of my evidence set. At least some of the elements of my evidence must be things other than objects. Our first proposal is false, and it cannot be fixed no matter how we specify the relation of acquaintance.

So now let's consider a second proposal:

S's evidence set at t is the set of all and only those property instantiations (i.e., states) and relations with which S is acquainted at t .

Once again, putting aside worries about the nature of acquaintance, can this proposal avoid the problem encountered by our first proposal? No. Suppose my evidence set includes not the knife, but rather the state of the *knife's being bloody*. My evidence includes not the kitchen, but rather the relation of *the knife's being in the kitchen*. Will having such states and relations in my evidence set now fully determine (to the extent that it is determinate) how confident it is rational for me to be in various hypotheses? Not unless it is impossible for there to be a pair of situations that differ with respect to how confident I can rationally be, but that are the same with respect to which states and relations I am acquainted with, and so with respect to what is (on the present proposal) in my evidence set. But this is not impossible. For instance, suppose that the state of the knife's being bloody is identical to the state of knife's being covered in a fluid that has the chemical composition C (where C is, say, the chemical composition distinctive of blood). But of course it's possible for me to have no

reason to accept the true identity statement that the state of being bloody is identical to the state of being covered in a fluid with chemical composition C. If I have no reason to accept this true identity statement, then, even if I am acquainted with the state of the knife's being covered in a fluid that has the chemical composition C, that still doesn't determine how confident I should be about the defendant's guilt. This example is representative. My evidence set, therefore, cannot consist exclusively of such states.

The same considerations that apply to states also apply to relations: relations, no more than states, determine—as precisely as it is determined—which propositions I rationally ought to be confident of. This is because, once again, it is possible for there to be a pair of situations that differ with respect to what I can rationally be confident of, but that are the same with respect to which states and relations I am acquainted with. And notice, nothing about my argument for this conclusion has rested on any particular construal of the notion of acquaintance.

I conclude that our evidence set cannot consist exclusively of objects, states, or relations. Our evidence set must also contain some item(s) *Z* such that it is impossible for there to be a pair of situations that differ with respect to what I can rationally be confident of, but that are the same with respect to *Z*'s. But what could such *Z*'s be, if they are not objects, states, or relations? So far as I can see, the only things that are not objects, states, or relations, and are even remotely plausible candidates to be elements of our evidence set, are propositions of some kind or other. Perhaps the *Z*'s in question are true propositions. Perhaps they are known propositions. Perhaps they are believed propositions. But whatever they are, they are not objects, states or relations, and so, for lack of any other remotely plausible candidates, I conclude that they are propositions of some kind or other. One's evidence set includes some propositions.

Of course, this argument—besides having the disappointing defeasibility characteristic of all 'what else could it be?' arguments—also leaves it open that an evidence set has some members that are propositions, and other members that are not propositions. I will not attempt to argue against that hypothesis in this section. Rather, I will argue against it only later, once I have put forward my own account of evidence. For now, I will simply assume that this hypothesis is false. That is to say, I will assume that, if some of the elements of one's evidence set are propositions, then all of them are propositions. The only consideration that I can cite in defense of this assumption right now (prior to offering my own positive account of evidence) are methodological: other things being equal, an account of evidence according to which the elements of our evidence set are all of the same ontological category is to be preferred to any other account of evidence. And so, on the basis of this methodological point, in tandem with the

arguments above concerning objects, states, and relations, I conclude that all the elements of one's evidence set are propositions. All evidence is propositional.

Now the next question we should address is: which kinds of propositions are the elements of one's evidence set? That is, what properties must a proposition have in order to be an element of someone's evidence set at a given time? In the next section, I will canvass various possible answers to this question. Then, in the following section, I will state and defend my own answer to this question.

3 Which kinds of propositions are in one's evidence set?

We can divide up propositional accounts of evidence into two mutually exclusive and jointly exhaustive categories: the doxastic and the non-doxastic. Doxastic accounts of evidence entail that it is a necessary condition of p being an element of S's evidence set at time t that S believes that p at t . Non-doxastic accounts do not entail this. In this section, I will first critically examine various doxastic accounts, and then I will critically examine various non-doxastic accounts. I will end up accepting a non-doxastic account. But first, let's see what's wrong with various doxastic accounts. Of course, I cannot examine all possible doxastic accounts. What I will do instead is to examine a great variety of possible doxastic accounts: some that are somewhat plausible, and some that are not at all plausible. I will locate a problem of a certain sort with each such account. I hope that this procedure, tedious as it may be (especially in the attention it devotes to some implausible accounts of evidence), suggests a pattern to the reader, and that this survey thereby helps to reveal a systematic defect with all doxastic theories—not simply the particular theories that I survey.

3.1 Doxastic accounts of evidence

Doxastic account 1. p is a member of S's evidence set at t if and only S believes that p at t . For short, let's call this the 'E = B' account of evidence.

Counterexample. Suppose I am drawing billiard balls out of an urn one by one.² I know that there have been exactly 99 drawings so far, but I have not been watching the results, so I have no reason to believe anything very specific about what colors the drawn billiard balls have been so far. If I believe, for no good reason at all, that the first drawing produced a black ball, and that the second drawing produced a black ball, and that . . . , and that the 99th drawing produced a black ball, then, according to E = B, my evidence set includes all of those 99 propositions, along with the proposition that there have been exactly 99 drawings. Now, if, as E = B predicts, my evidence does include all of those propositions, and if furthermore I have no evidence against the proposition that the next ball will be black, then I can rationally be quite

² Williamson ([2000], Ch. 9) suggests a counterexample like the present one to E = B.

confident that the next ball will be black. Nonetheless, it is clear that, a situation of the sort that I have described—one in which I have not been looking at the drawn balls, and I have no reason to believe anything very specific about their colors—I cannot rationally be confident that the next ball will be black. Thus, the prediction issued by $E = B$ for the situation envisaged is false. And so $E = B$ is false.

Doxastic account 2. p is a member of S 's evidence set at t if and only if S believes that p at t , and p is true. Let's call this the ' $E = TB$ ' account of evidence.

Counterexample. Consider a scenario just like the one given as a counterexample to $E = B$, except that, as it happens (entirely unbeknownst to me), it is true that each of the first 99 drawings produced a black ball. Still, in such a situation, I cannot rationally be confident that the next ball will (even likely) be black. And yet $E = TB$ predicts that I can rationally be confident of this, because it says that I have no evidence against the next ball's being black, and I have plenty of evidence for it. After all, according to $E = TB$, my evidence consists of the following believed truths: ball #1 was black, ball #2 was black, . . . , ball #99 was black, and there were no other balls drawn. And so $E = TB$ issues a false prediction about the envisaged situation. $E = TB$ is false.

Doxastic account 3. p is a member of S 's evidence set at t if and only if S justifiably believes that p at t . Let's call this the ' $E = JB$ ' account of evidence.

Counterexample. If I know that the first 99 drawings produced all black balls, then I am justified in believing that the next drawing will produce a black ball. But according to $E = JB$, that latter proposition gets to go into my evidence set. So now my evidence includes that the first 100 drawings produced all black balls, and so I am justified in believing that the next drawing will produce a black ball. But according to $E = JB$, that latter proposition gets to go into my evidence set. So now my evidence set includes that the first 101 drawings produced all black balls, and so on. There is no end to this regress. $E = JB$ is false.

Note that the regress argument just given against $E = JB$ works equally well against another doxastic theory: p is a member of S 's evidence set at t if and only if S 's belief that p is highly confirmed at t . Again, if I know that the first 99 drawings produced all black balls, then my belief that the 100th drawing will produce a black ball is highly confirmed, and so my belief that the 101st drawing will produce a black ball is highly confirmed, and so my belief that the 102nd drawing will produce a black ball is highly confirmed, and so on.

Doxastic account 4. p is a member of S 's evidence set at t if and only if S 's belief that p is justified and true. Let's call this the ' $E = JTB$ ' account of evidence.

Counterexample. If I know that the first 99 drawings produced all black balls, then I am justified in believing that the 100th drawing will produce a black ball. And suppose it is true that the 100th drawing will produce a black ball. In that

case, according to $E = \text{JTB}$, that latter proposition gets to go into my evidence set. So now my evidence includes that the first 100 drawings produced all black balls, and so I am justified in believing that the 101st drawing will produce a black ball. Suppose again that it is true that the 101st drawing will produce a black ball. In that case, according to $E = \text{JTB}$, that latter proposition gets to go into my evidence set. So now my evidence set includes that the first 101 drawings produced all black balls, and so on. Once again, as long as it's true that the next drawing will produce a black ball, there is no end to this regress. $E = \text{JTB}$ is false.

Doxastic account 5. p is a member of S 's evidence set at t if and only if S 's belief that p is reliably formed or sustained at t . Let's call this the ' $E = \text{RB}$ ' account of evidence.

Counterexample. Let N be a complicated neurological state which, as a matter of fact, is such that S is in N when and only when S believes that she is in N . Furthermore, let's suppose that S is entirely unaware of the truth of this biconditional: S is simply surprised occasionally to find herself believing that she is in N . (She might find herself believing this in something like the way that many of us, when we are in a strange, new place, find ourselves believing—for no apparent reason—that we have seen this strange place before, somewhere.) In that case, S 's beliefs that she is in N will be as reliably formed or sustained as any belief could be. But rationality does not require that S distribute her confidence over hypotheses in proportion to the degree of support that those hypotheses receive from her belief—perplexing as it is by her own lights—that she is in N . Thus, $E = \text{RB}$ is false.

Notice that the same counterexample just given against $E = \text{RB}$ would work equally well against the view that p is a member of S 's evidence set at t if and only if S 's belief that p is reliably formed or sustained at t , and is also true.

Doxastic account 6. p is a member of S 's evidence set at t if and only if S 's belief that p is blamelessly held at t . Let's call this the ' $E = \text{BB}$ ' account of evidence.

Counterexample. Through no fault of her own, S uncritically believes everything that she hears other people assert, at least when those assertions go uncontradicted by other assertions. She cannot be blamed for doing so, because she is somehow prevented from exercising her critical capacities very well, and indeed, we may suppose that her beliefs end up being right more often when she does not exercise her critical capacities, and she knows this. So she blamelessly imbibes the following beliefs from her fellows: all of the books in the local library are full of lies and corrupt the mind of anyone who reads them, scientific inquiry is a lot of highly funded academic babble with no better epistemic credentials than astrology, and foreigners are all out to kill us. It is, of course, a terrible shame that S 's cognitive powers prevent her from effectively examining these beliefs, but that is not S 's fault. Clearly, rationality does not require that

S distribute her confidence across hypotheses in proportion to the degree that those hypotheses are supported by the blamelessly held beliefs just mentioned. $E = BB$ is therefore false.

Notice that the same counterexample that refutes $E = BB$ also refutes the view that p is a member of S's evidence set at t if and only if S's belief that p is both blameless and true at t .

Doxastic account 7. p is a member of S's evidence set at t if and only if p is a member of the largest coherent subset of S's beliefs at t .

Counterexample. Consider Alvin Plantinga's Case of the Epistemically Inflexible Climber.

'Ric is climbing Guide's Wall, on Storm Point in the Grand Tetons; having just led the difficult next to last pitch, he is seated on a comfortable ledge. [...] He believes that Cascade Canyon is down to his left, that the cliffs of Mount Owen are directly in front of him, that there is a hawk gliding in lazy circles 200 feet below him, that he is wearing his new *Fire* rock shoes, and so on. His beliefs, we may stipulate, are coherent. Now add that Ric is struck by a wayward burst of high-energy cosmic radiation. This induces a cognitive malfunction; his beliefs become fixed, no longer responsive to changes in experience. No matter what his experience, his beliefs remain the same. At the cost of considerable effort his partner gets him down and, in a desperate last-ditch attempt at therapy, takes him to the opera in nearby Jackson, where the New York Metropolitan Opera on tour is performing *La Traviata*. Ric is appeared to in the same way as everyone else there; he is inundated by wave after wave of golden sound. Sadly enough, the effort at therapy fails; Ric's beliefs remain fixed and wholly unresponsive to his experience; he still believes that he is on the belay ledge at the top of the next to last pitch of Guide's Wall, that Cascade Canyon is down to his left, that there is a hawk sailing in lazy circles 200 feet below him, that he is wearing his new *Fire* rock shoes, and so on.'

(Plantinga [1993], p. 82)

Both before and after Ric suffers from the cognitive malfunction induced by the burst of cosmic radiation, he has the same largest coherent subset of beliefs. But, after he suffers from the cognitive malfunction, and his beliefs are radically disjoint from his experience, rationality does not require that Ric distribute his confidence across hypotheses in proportion to the degree to which those hypotheses are supported by the beliefs in that largest coherent set. Thus, after he suffers from the cognitive malfunction, the beliefs in that largest coherent set cannot constitute Ric's evidence set.³

³ Some philosophers may worry that Plantinga's thought experiment is not coherent: we cannot coherently stipulate that beliefs stay fixed while experiences change. It may help to allay this worry to alter the thought experiment slightly: instead of stipulating that Ric's beliefs stay completely fixed, let's allow that a few of his beliefs change with his experiences, but the large preponderance of his beliefs stay fixed. We can run the same argument with the thought experiment so emended.

Doxastic account 8. p is a member of S 's evidence set at t if and only if S knows that p at t . Following Williamson, let's call this the 'E = K' account of evidence.

Counterexample. Imagine two subjects who differ slightly in their phenomenal states. One of them is enjoying a perfectly clear visual image of a 4-speckled hen against an otherwise simple and uniform background, and the other is enjoying an equally clear and otherwise identical visual image of a 5-speckled hen. Suppose furthermore that neither of them believes or knows anything very specific about the number of speckles on the hen image that they're now having, i.e., neither have any beliefs as to whether the number of speckles is four or five, though each (we may suppose) believes that the number of speckles is greater than three and less than seven. The two subjects differ, therefore, both in their phenomenal states, but—let's suppose—not in any other way that is independent of this stipulated differences in their phenomenal states. Now, let's consider two issues: first, should each subject distribute her credences in the very same way over the hypothesis that she has an image of a 4-speckled hen and the hypothesis that she has an image of a 5-speckled? And second, if the two subjects should distribute their credences differently over those two hypotheses, is this difference explicable by appeal to the difference in the propositions that they know to be true?

To answer the first question, let's first consider how confident the 4-speckle subject should be that the visualized hen has four speckles. We may assume that the speckles are spaced apart in such a way as to be obviously distinct from each other, but also to be surveyable all at once. Now, if the subject has some reason to believe that she is confused, or that she is very bad at counting speckles, then perhaps she should not be very confident that the visualized hen has four speckles. But if the subject has no reason to believe that she suffers from either of these defects (and we may coherently suppose that she does not have any such reason), then she ought to be very confident that the hen image has four speckles. The corresponding point holds for the subject whose visualized hen has five speckles: she ought to be very confident that the hen image has five speckles. Thus, the two subjects ought to distribute their confidence very differently from each other across the two hypotheses: (i) the hen image has four speckles, and (ii) the hen image has five speckles.

I anticipate the following objection: if the 4-speckle subject does not know that the hen image has four speckles, then she must be somehow confused or distracted. But if she is confused or distracted, she is not *able* to survey the four speckles accurately. And if she is not able to survey them accurately, then rationality cannot require that she be confident that there are four speckles, for rationality cannot make demands of us that we cannot comply with.

A problem with this objection is that the inference from the subject's confusion or distraction to her inability to survey the four speckles accurately is not

a good inference. I may confusedly or distractedly call my spouse by someone else's name, but this does not imply that I am not able to call my spouse by her correct name, let alone that rationality does not require me to call my spouse by her correct name. I may confusedly or distractedly run the red light, but this does not imply that I am not able to stop at the red light, let alone that rationality does not require me to stop at the red light. Therefore, the present objection fails. I conclude that rationality requires the two subjects to distribute their confidence differently across the 4-speckle hypothesis and the 5-speckle hypothesis.

But, now turning to the second of our two questions above, can this difference be explained by appeal to the difference in which propositions the two subjects know to be true? It may seem obvious that it cannot be so explained; does it not follow from the stipulations of our example that the two subjects know all the same propositions? No. What was stipulated was that the two subjects are identical in all respects that are *independent of the specified phenomenal difference between them*. But it may be that the specified phenomenal difference cannot obtain between the two subjects unless there is also a difference between what one knows and what the other knows. Specifically, since the two subjects differ in their phenomenal states—one of them enjoying a 4-speckled hen image while the other enjoys a 5-speckled hen image—it may be thought that they must also differ in their knowledge, for each subject knows *de re*, of each speckle in her image, that *it* (attending to that speckle) is *there* (attending to its position in the image). Thus, the 4-speckle subject will have *de re* knowledge of four such propositions (one for each speckle in her image), whereas the 5-speckle subject will have *de re* knowledge of five such propositions. Now, I do not know if the stipulated phenomenal difference does require such an epistemic difference or not: if it doesn't require such an epistemic difference, then the two subjects are rationally required to distribute their confidence differently, even though they know all the same propositions (adjusting for differences in indexicals). In that case, $E = K$ is false.

So let's suppose, for the sake of argument, that the stipulated phenomenal difference does require such an epistemic difference, i.e., the two subjects cannot differ with respect to how many speckles are in their image unless they also differ with respect to which propositions of the form

that (attending to a particular speckle) is *there* (attending to a particular spot in one's visual field)

they each know to be true. I will now argue that, even if the subjects do differ in which propositions of this form they know to be true, that latter difference cannot explain why they differ in their rational distribution of confidence over the 4-speckle and the 5-speckle hypotheses.

Consider the 5-speckle subject: she knows each of the following five propositions to be true:

That (attending to the first speckle) is *there* (attending to its location)

That (attending to the second speckle) is *there* (attending to its location)

That (attending to the third speckle) is *there* (attending to its location)

That (attending to the fourth speckle) is *there* (attending to its location)

That (attending to the fifth speckle) is *there* (attending to its location)

I'll refer to such truths as 'speckle truths,' and I'll say that the five speckle truths above 'belong to' the 5-speckle subject. (In general, I'll say that a speckle truth 'belongs to' a subject just in case it is a truth about a particular speckle in that subject's visualized hen. Alternatively, I'll say that the speckle truth in question is 'her own.')

The 5-speckle subject knows five speckle truths that belong to her. The 4-speckle subject knows the four speckle truths that belong to her. Does this difference between the sets of truths that they know help to explain why the 5-speckle subject should have greater confidence in the 5-speckle hypothesis than the 4-speckle subject should, or why the 4-speckle subject should have greater confidence in the 4-speckle hypothesis than the 5-speckle subject should?

No, it cannot. To see why this is the case, consider a situation in which the 4-speckle subject knows all of the following five speckle truths that belong to her:

That (attending to the first speckle) is *there* (attending to its location)

That (attending to the second speckle) is *there* (attending to its location)

That (attending to the third speckle) is *there* (attending to its location)

That (attending to the fourth speckle, by attending leftward from the third speckle) is *there* (attending to its location)

That (attending to the fourth speckle, by attending upward from the first speckle) is *there* (attending to its location)

Even if the 4-speckle subject knows each of the five speckle truths listed above, it is still the case that the 4-speckle subject should be more confident of the 4-speckle hypothesis than of the 5-speckle hypothesis. But in a situation like the one just described, the 4-speckle subject knows just as many of her own speckle truths as our envisioned 5-speckle subject does: namely, five. Since they know the same number of speckle truths, and yet they still should distribute their confidence differently between the 4-speckle and 5-speckle hypotheses, the explanation of this latter difference cannot appeal to the fact (even when it is a fact) that they differ in how many speckle truths they know. The difference in how many speckle truths they know could disappear, even though they should still distribute their confidence differently between the 4-speckle and the 5-speckle hypotheses.

Finally, it might be proposed that, even when the 4-speckle subject and the 5-speckle subject know the same number of speckle truths, what explains the

difference in how they should distribute their confidence is not anything about the *number* of speckle truths that they know, but is rather something about which particular speckle truths they know. But what could this latter difference be, such that it explains the difference in how they should each distribute their confidence? I do not see any way for the defender of $E = K$ to answer this question plausibly. And so I conclude that, even if our hen-visualizing subjects do not know all the same truths, the difference in which truths they know still cannot explain the difference in how confident it is rational for them to be in, say, the 4-speckle hypothesis. But the difference in their evidence must explain the difference in how confident it is rational for them to be in the 4-speckle hypothesis. And thus $E = K$ is false: in general, propositions that a subject knows to be true cannot play the explanatory role that the subject's evidence is supposed to play.

While the foregoing survey of eight doxastic accounts of evidence cannot come close to exhaustiveness, it strongly suggests that doxastic accounts of evidence are in trouble. And perhaps the basis for rejecting $E = K$ can serve as a basis for rejecting them all: some of the propositions in our evidence set (e.g., the proposition that I am having a 4-speckled hen image) are propositions that get to be in our evidence set even though we do not believe them. They get to be in our evidence set by virtue of something else. So what is this other non-doxastic factor? Let's now turn our attention to that question, and survey some non-doxastic accounts of evidence.

3.2 Non-doxastic accounts of evidence

There is a difference between *there being evidence* for a certain hypothesis and *your having that evidence*. The existence of evidence does not suffice for you to *have* that evidence. But what is involved then, in your having that evidence? At least this much: there must be some kind of epistemic relation (using the term 'epistemic' here to encompass all the various properties studied in epistemology, and not simply the property of knowledge) between you and the evidence propositions. As we saw in the preceding section, the epistemic relation in question need not involve or require the relation of belief: you need not believe the propositions that are in your evidence set. But you must somehow or other be epistemically related to them, or else you don't have the evidence. And you must be epistemically related to those propositions at a particular time, or else you don't have the evidence in question at that time. What kind of epistemic relation is at issue here? Different non-doxastic accounts of evidence offer different answers to this question.

Non-doxastic account 1. *p* is a member of S's evidence set at *t* if and only if, at *t*, S has a propositional attitude towards the proposition *p*.

Counterexample. John wants it to be the case that Kerry wins the election. But rationality does not require that John distribute his confidence across hypotheses in proportion to the degree to which those hypotheses are supported by a conjunction one conjunct of which is the proposition that Kerry wins the election: the proposition that Kerry wins the election should not help to fix John's degrees of confidence. One problem with non-doxastic account 1 is that wanting it to be the case that p does not, by itself, give me any reason to believe that p , or to believe anything that is supported by p .

Notice that a trivial variation of the preceding counterexample to non-doxastic account 1 also serves to refute the view that p is a member of S's evidence set at t if and only if, at t , S has a propositional attitude towards the proposition p , and *furthermore* p is true. (Replace 'Kerry wins the election' with 'Bush wins the election.')

Non-doxastic account 2. p is a member of S's evidence set at t if and only if, at t , S has a non-motivating propositional attitude towards the proposition p . (A 'non-motivating' propositional attitude is one that does not provide motivation to do anything. For present purposes, I simply assume the truth of the common Humean view that the category of 'non-motivating' propositional attitudes excludes desires, hopes, and intentions, but includes, *inter alia*, beliefs.)

Counterexample. In order to construct a *reductio ad absurdum* proof of p , S begins by assuming not- p . Nonetheless, rationality does not require that S distribute her confidence in hypotheses in proportion to the degree to which they are supported by a conjunction, one conjunct of which is not- p : not- p is not a member of S's evidence set, despite S's supposing that not- p . But supposition is a non-motivating propositional attitude, so not all of the propositions towards which S has a non-motivating propositional attitude at t are in S's evidence set at t . The problem with non-doxastic account 2 is that not all non-motivating propositional attitudes provide reasons to believe their propositional contents. We need to restrict the class of relevant propositional attitudes even more in order to arrive at a possibly correct account of evidence.

Notice that a trivial variation of the preceding counterexample to non-doxastic account 2 also serves to refute the view that p is a member of S's evidence set at t if and only if, at t , S has a non-motivating propositional attitude towards the proposition p , and *furthermore* p is true.

Non-doxastic account 3. p is a member of S's evidence set at t if and only if, at t , S has a conviction-carrying attitude that has the representational content that p . (A 'conviction-carrying' propositional attitude that p is an attitude that either generally involves, or has a tendency to produce, the conviction that p . Thus, motivating propositional attitudes are none of them conviction-carrying, and neither is supposition conviction-carrying, but believing, perceiving, and remembering are all conviction-carrying.)

Counterexample. If I believe, for no reason whatever, that the first 99 drawings of a ball from a particular urn have produced black balls, rationality does not thereby require that I am confident that the 100th drawing will produce a black ball. Not all of my beliefs are in my evidence set, and so not all of the propositions towards which I bear conviction-carrying, non-motivating propositional attitudes are in my evidence set. So non-doxastic account 3 provides too generous an account of my evidence.

Notice that the same counterexample to non-doxastic account 3 can serve to refute the view that p is a member of S's evidence set at t if and only if, at t , S has a conviction-carrying propositional attitude that has the representational content that p , and p is true. We need only stipulate that, entirely unbeknownst to me, the first 99 drawings did indeed all produce black balls.

Non-doxastic account 4. p is a member of S's evidence set at t if and only if, at t , S has a conviction-producing *but not conviction-involving* propositional attitude that has the representational content that p . (To say that an attitude is not conviction-involving is to say that it is not a necessary condition of having such an attitude with respect to the proposition p that one have the conviction that p . Some conviction-producing attitudes are not conviction-involving: for instance, *seeming to see that p* is conviction-producing, but not conviction-involving, and so is *seeming to remember that p* .)

Counterexample. Suppose you learn about a peculiar psychosis that invariably causes its victims to not believe that they exist. You are a hypochondriac with a tendency to worry that you have every disorder that you learn about, so in this case you attempt to ascertain that you do not have this disorder. This attempt succeeds easily because you have conclusive evidence that you do not suffer from the disorder, namely, that you do believe that you exist. But, while the proposition that you believe that you exist is the content of some of your conviction-carrying attitudes (e.g., you believe that you believe that you exist, you know that you believe that you exist, you understand that you believe that you exist, perhaps you introspect that you believe that you exist) it is not a necessary condition of that proposition being in your evidence set that it be the content of any attitude that is not conviction-involving. Thus, what is obviously a bit of your evidence does not get to count as an element of your evidence set by non-doxastic account 4, and so non-doxastic account 4 is false.

Notice again that our counterexample, as it stands, works equally well to refute the view that p is a member of S's evidence set at t if and only if, at t , S has a conviction-producing but not conviction-involving propositional attitude that has the representational content that p , and p is true. It also works equally well to refute the view that p is a member of S's evidence set at t if and only if, at t , S has a conviction-producing but not conviction-involving propositional attitude that has the representational content that p , and *one's having this attitude is reliably caused by the fact that p* .

Does the preceding counterexample refute the view that p is a member of S 's evidence set at t if and only if, at t , S has a perceptual experience or apparent memory that has the representational content that p ? That's not clear. It depends upon whether or not one's evidence can include the proposition that one believes that one exists without one's having a perceptual experience or apparent memory with the representational content that one believes that one exists. And that issue is not settled.

So it seems that we *might* finally have a counterexample-free non-doxastic account of evidence as follows: the proposition that p is a member of S 's evidence set at t if and only if, at t , S has a perceptual experience or apparent memory that has the representational content that p .

Alternatively, here's another non-doxastic account that's not refutable by means of any of the considerations adduced so far: p is a member of S 's evidence set at t if and only if p is a true proposition to the effect that, at t , S has a particular perceptual experience or apparent memory.

Why should we not settle for one or the other of the two accounts just suggested? The worry I have about each of these two accounts is that neither of them is sufficiently well motivated. Maybe it is true that p is a member of S 's evidence set at t if and only if, at t , S has a perceptual experience or apparent memory that has the representational content that p . Or, maybe it is true that p is a member of S 's evidence set at t if and only if p is a true proposition to the effect that, at t , S has a particular perceptual experience or apparent memory. In fact, each of these two views is compatible with the view that I will defend below, and neither is entirely implausible. But what's so special about perceptual experiences or apparent memories? Why is it that they get to determine our rational degrees of confidence, but other mental states do not? We should need to hear a good answer to this question in order to be moved to accept either of these two non-doxastic account of evidence. Until I hear such an answer, I leave it open that either account is correct, but I regard each as insufficiently well motivated: they leave us in the dark as to why it is that just these mental states determine what's in our evidence set, and so determine how we should distribute our confidence in hypotheses.

We could, of course, continue to generate possible non-doxastic accounts of evidence, and check to see if they are counterexample-free *and* sufficiently well motivated. But our search so far has yielded only negative results, and so we might, by now, have gotten tired of this trial-and-error procedure. To the fatigued inquirer, I suggest an alternative course of action. Let's design a non-doxastic account of evidence that's guaranteed to satisfy both of our two conditions of adequacy on such an account, and let's do this by *writing those conditions into the account itself* (and then defending the resulting account against objections).

Some philosophers may protest that what we produce by following this procedure is not an *account* of evidence. Nor does it satisfy our two conditions of adequacy, and that's because it does not give any informative *explanation* of why evidence can be used in the rational regulation of our attitudes, or of why the requirement of total evidence is true. That is, although we may obtain explanations of these things by simply writing them into the account, the explanations that we thereby obtain are not informative. We are simply saying the following: what it is for a subject to have a certain evidence set E is just for a subject to be able to use all and only the elements of E in the rational regulation of her attitudes, and for rationality to require that that subject distribute her confidence over hypotheses in proportion to the degree to which those hypotheses are supported by the conjunction of all and only the elements of E. Because this 'account' does not offer an informative explanation of these features of having evidence, it is perhaps more aptly described as an *idea* about what it is to have evidence, rather than an *account* of evidence: it does not really *account* for the features of evidence mentioned in our two conditions of adequacy, except, perhaps, by suggesting that they are fundamental features of evidence. (To say that they are fundamental features of evidence is to say that evidence, simply by its nature, has these features, and furthermore there is no further explanation of its having these features. It is, perhaps, a fundamental feature of water that it is H₂O, for there is no further explanation of what makes it the case that water is H₂O, besides the uninteresting claim, 'that's just what it is!'.)

Furthermore, this proposed idea about what it is to have evidence—the idea that I advance here—might seem, to many philosophers, to enjoy all the virtues of theft over honest toil. And such philosophers are right: my proposed idea does indeed enjoy all the considerable virtues of theft over honest toil. Consequently, I call it the 'larcenous idea of evidence,' or the LIE for short.

The LIE will need some elaboration and defense. At the very least, we must show how it is that we can produce a coherent and workable idea of evidence by making use of our two conditions of adequacy. And we should also see what the consequences of the resulting idea are, and what they are not.

4 Elaborating and defending the LIE

At the end of the preceding section, I described the strategy that I will employ now: rather than constructing an independently motivated account of evidence that can be used to predict and explain the two conditions of adequacy imposed at the beginning of this paper, I will instead simply appeal to those two conditions in order to construct a particular view about what it is for a proposition to be in a subject's evidence set at a particular time.

First, by appeal to the requirement of total evidence, we can say the following: a person *S*'s total evidence set *E* at a time *t* contains just those propositions the conjunction of which, call it ' E_{conj} ,' is such that rationality requires *S* to distribute her confidence at *t* across hypotheses in proportion to the support that those hypotheses receive from E_{conj} . While this imposes a substantial constraint on *S*'s evidence at *t*, it does not impose a constraint sufficient to determine uniquely the elements of *S*'s evidence set. That is because, for any proposition, there are many different conjunctions that are logically equivalent to that same proposition (or, as I shall say, many different ways of 'decomposing' that proposition into conjuncts), and some decompositions of E_{conj} into conjuncts are such that, obviously, not every conjunct in the decomposition is in that person's evidence set. For instance, it is consistent with the constraint just stated that *S*'s evidence set at *t* has just the following two members:

- (i) the disjunction: E_{conj} or Napoleon had an even number of eyelashes when he was born
- (ii) the disjunction: E_{conj} or Napoleon had an odd number of eyelashes when he was born.

No matter what E_{conj} is, it is going to be logically equivalent to the conjunction of (i) and (ii). And so E_{conj} (no matter what it is) can be decomposed into conjuncts (i) and (ii). But no actual person's evidence set currently contains either (i) or (ii): no one can use (i), say, in the rational regulation of her attitudes. For instance, no one can rationally boost their confidence in the hypothesis that, if E_{conj} is false then Napoleon had an even number of eyelashes when he was born on the grounds that that latter hypothesis follows from (i). And so this first constraint on a person's evidence set does not suffice to determine which propositions are elements of that person's evidence set.

So let's supplement this first constraint (taken from our second condition of adequacy) with a second constraint (taken from our first condition of adequacy), namely: the correct decomposition of E_{conj} into conjuncts is a decomposition which is such that, for every conjunct in that decomposition, *S* is able to use that conjunct in the rational regulation of her attitudes. If there is more than one decomposition of E_{conj} that satisfies this second constraint, that may be fine: we may let all of the conjuncts in every such decomposition be elements of *S*'s evidence set: their conjunction will still be E_{conj} , and so will still satisfy our first constraint.

But there is a worry: what if there is no decomposition of E_{conj} that satisfies this second constraint? And so what if there is no set of propositions that satisfies both of the two constraints that we're using to specify the elements of *S*'s evidence set? Is that possible? No, it is not.

To see why this is not possible, let's recall that E_{conj} is a proposition such that rationally requires me to distribute my confidence across hypotheses in

proportion to the support that those hypotheses receive from E_{conj} . Now, if this is what rationality requires of me, then I must have the ability to comply with this requirement: governments, armies, and social conventions may require me to do all sorts of things that I lack the ability to do, but rationality itself cannot require me to do anything that I lack the ability to do.

Let me be clear that, to say that S has the ability to do something is not to say that it is metaphysically possible for S to do it. In general, S's having the ability to F does not imply that there is a possible situation in which S F's. By virtue of knowing how to add, we have the ability to compute infinitely many different sums. But this doesn't imply that there is a possible world or a possible situation in which we do compute infinitely many different sums. It doesn't even imply that, for any given sum, there is a possible world or a possible situation in which we compute it: some sums may take longer than any possible human life span to compute, given the computational limits of our brain, and the biological limits on the longevity of our heart, liver, and brain. In general, abilities are not metaphysical possibilities: or, to put the point in terms that Chomsky made popular, performance may be restricted in ways that competence is not restricted. But in order for norms of rationality to apply to us, we must have the abilities (or competences) necessary to comply with those norms. And specifically, in order for us to be rationally required to do something, we must have the ability to do it, whether or not it is metaphysically possible for us to do it. So it's a necessary condition of S's being rationally required to proportion her confidence across hypotheses in proportion to the support that those hypotheses receive from E_{conj} that S has the ability to do this.

I must, therefore, have the ability to distribute my confidence across hypotheses in proportion to the support that those hypotheses receive from E_{conj} . But then—and this is the crucial point—I must have the ability to use E_{conj} itself in the rational regulation of my attitudes. And so E_{conj} itself is an element of my evidence set. And this implies that there is guaranteed to be at least one decomposition of E_{conj} into conjuncts that satisfies our second constraint: namely, a decomposition that includes just E_{conj} itself. There may, of course, be many other decompositions of E_{conj} into conjuncts that satisfy our second constraint. But there is guaranteed to be at least one. And that is all I need to show that our two conditions of adequacy on an account of evidence are jointly satisfiable.

I have just said that we can use E_{conj} itself—the conjunction of all of our evidence—in the rational regulation of our attitudes. But this statement is bound to seem puzzling to many people—especially to those who engage in, or study, scientific practice. In the course of scientific practice, we typically regulate our attitudes in the light of new evidence in a more or less *piecemeal* fashion. How can we possibly use our *total* evidence, E_{conj} itself, in the rational regulation of our attitudes? Is there a concrete example of our doing this?

The clearest examples of our using E_{conj} itself in the rational regulation of our attitudes are cases that do not come up in the course of scientific practice, but rather only come up when we have to rule out utterly preposterous challenges to the whole framework in which science is conducted. For instance, when I correspond with my collaborators, or when I read *Nature*, or when I ask my research assistant about the results of the experiment that I asked her to run, I firmly believe that the testimony that I receive from these diverse sources is not part of an enormous and thoroughgoing conspiracy to deceive me. If someone were to challenge me to defend my belief that this testimony is not all part of an enormous and thoroughgoing conspiracy to deceive me, the only way in which I could meet the challenge would be to rely upon my total evidence and reply ‘there is *no indication whatsoever* that it is a conspiracy to deceive me, and *every indication* that it is not!’

In order for me to do science at all, I must confidently accept a great deal of the testimony that I receive, and so I must be confident that such testimony is, at least by and large, not the product of deception. What makes me rational in being confident that such testimony is, at least by and large, not the product of deception, cannot be *merely* the evidence that I acquire in the course of doing science, for when I began my scientific career I was already rationally confident that textbooks and teachers and so on were not all engaged in a massive conspiracy to deceive me. Rather, it is my *total* evidence, E_{conj} , that makes me rational in being confident that such testimony is, at least by and large, correct.⁴

Can I spell out the content of E_{conj} ? Can I tell you what all of my evidence is? The only way that I know how to tell you is by ostension: *this*—I might say (perhaps with a vague sweeping gesture)—is all my evidence. Whatever it is that I am ostending by means of this demonstrative, it is what makes it rational for me to be confident that the testimony that I receive in the course of doing science is not part of an enormous conspiracy to deceive me.⁵ In general, the

⁴ There are two popular alternatives to claiming that my total evidence makes it rational for me to be confident that I am not the victim of massive deception. According to the first alternative, it is simply *not* rational for me to be confident of this (my confidence is not open to assessment as rational or not). According to the second alternative, it is *a priori* rational for me to be confident of this. The first option is ruled out by the fact that what it is rational for me to believe is closed under known entailment. I argue against the second option in my ([2007]).

⁵ Our procedure of using our total evidence in the rational regulation of our attitudes is so commonplace as to be nearly invisible. But it was just this procedure that G. E. Moore illustrated vividly in his *Proof of an External World*. For instance, on the basis of all of the evidence that I now have, I am now rationally confident that I have two hands, that I am a human being living on the surface of the planet Earth, speaking English, and teaching at a university. My total evidence is relevant to my confidence in each of these propositions, and that’s because my confidence in each of these propositions is largely based on my confidence that this is not all a dream, and my confidence that this is not all a dream is in turn based on my total evidence. This is, I believe, why Moore says that in order to prove that here are two hands, he would need to provide *all* of his evidence. See (Moore [1993]).

ordinary cases in which we use E_{conj} in the rational regulation of our attitudes are cases that don't come up *within* the course of scientific practice, but are rather the kinds of cases that make it possible for us rationally to engage in scientific practice in the first place. This is why such cases are more commonly studied by epistemologists than by philosophers of science.

We are, then, guaranteed to find at least one (trivial) decomposition of S's total evidence into conjuncts, each of which S is able to use in the rational regulation of her attitudes, and what we've said so far is just the following: that a proposition is an element of S's evidence set if and only if it is a conjunct in some such decomposition.

To sum up our positive conclusions about evidence: p is an element of S's evidence set if and only if, if E_{conj} is a proposition which is such that S is rationally required to distribute her confidence across hypotheses in proportion to the support that those hypotheses receive from E_{conj} , then p is a member of some decomposition of E_{conj} into conjuncts, each of which S is able to use in the rational regulation of her attitudes. For any person who has evidence at all, there is guaranteed to be at least one such decomposition: namely, that which includes E_{conj} itself.

LIE is a general claim about the nature of membership in a subject's evidence set. By itself, it does not imply anything about which non-normative facts determine whether any particular proposition is a member of a subject's evidence set at a time. Thereby, it leaves plenty of room for controversy about how (or even whether) the non-normative facts fix one's evidence. Many philosophers will find cause for complaint in this feature of LIE. Specifically, they will complain that LIE cannot help us to address the basic normative issues that are at stake in the question of what the elements of one's evidence set are. The LIE cannot point us towards the non-normative considerations to which we can appeal in addressing these normative issues.⁶

In response to this complaint about the LIE, I want to concede right away: the LIE does not offer any way to settle, in non-normative terms, the normative question of what the elements of one's evidence set are. But I believe that any *fully general* idea of evidence *should not* attempt to settle that normative question in non-normative terms. Any attempt to settle this general normative question in non-normative terms would constitute an attempt to dictate to every form of inquiry—legal, scientific, or otherwise—what sorts of non-normative considerations determine what gets to be part of its evidence base for that form of inquiry. But, it seems to me, what non-normative considerations determine what gets to be part of the evidence available in particular form of inquiry is itself a matter to be worked out within that form of inquiry. Thus, astrophysicists

⁶ Why should we think that, in determining what evidence we have, we must appeal to non-normative considerations? Why can we not simply appeal to normative considerations?

have to figure out what non-normative facts determine the kind of thing that gets to be in the evidence base for astrophysics, clinical psychologists have to figure out what non-normative facts determine the kind of thing that gets to be in the evidence base for clinical psychology, ecologists have to figure out what non-normative facts determine the kind of thing that gets to be in the evidence base for ecology, and legislators and judges have to determine what non-normative facts determine the kind of thing that gets to be in the evidence base for various kinds of criminal or civil cases. Although it is remotely possible that there is some non-normative genus that subsumes as species all of these non-normative determinants of evidence, there is no reason to expect this to be the case, and the enormous apparent diversity of these disciplines gives us reason to expect it not to be the case.

Although the LIE is guaranteed to satisfy the two conditions of adequacy put forward at the beginning of this paper, it leaves a great many issues open. Here are some of the many general philosophical issues about evidence that are left open by the LIE:

1. Some philosophers (e.g., Joyce [2005]) wonder whether the evidentiary status of a proposition is a matter of degree: are some propositions *more* evidential for you at a particular time than others are? Are some propositions more evidential for you at a particular time than for me at that same time? The LIE, by itself, does not imply any answer to these questions.
2. Some empiricists say that all evidence propositions are directly observationally verifiable. But there is considerable disagreement both about what this constraint amounts to, and whether it is true in any form. The LIE, by itself, does not imply any answer to these questions.
3. Relativists (e.g., Rorty [1979]) say that which propositions get supported by evidence propositions is relative not just to a person and a time, but also to some further relatum—e.g., a disciplinary matrix, or a system of social norms, or a language-game, or what have you. Anti-relativists (e.g., Boghossian [2006]) deny this. The LIE, by itself, does not suggest any answer to these questions. Indeed, it implies nothing interesting about relations of evidential support.
4. Anti-intellectualists (e.g., Stanley [2005]) say that which propositions get to count as evidence propositions for a person at a time depends not just upon truth-conducive factors, but also upon that person's practical goals or interests. Again, LIE, by itself, does not suggest any position on this issue.

5. Contrastivists about evidence (e.g., Schaffer [2005]) say that the property of being in a specific person's evidence set at a specified time is not a monadic property, but is rather a relation that holds between an evidence proposition and a contrast proposition, and which of those various properties is denoted by a particular ascription of the form '*e* is in S's evidence set at *t*' is relative to a context of ascription. Others deny this claim. Again, LIE, by itself, does not suggest any position on this issue.
6. Evidentialists (e.g., Conee and Feldman [1985]) say that propositional justification supervenes solely upon what evidence one has: fix the evidence set that S has at time *t*, and it's thereby fixed what S has justification to believe at *t*. Anti-evidentialists (e.g., Fantl and McGrath [2002]) deny this. Again, LIE, by itself, does not suggest any position on this issue.

These are all substantial and important issues about the nature of evidence, but LIE does not address any of them. In these ways, among others, LIE falls far short of a complete account of what it is for a proposition to be in a subject's evidence set at a given time.

Acknowledgements

I am grateful to Bill Lycan, and also to the editors and an anonymous referee for *British Journal for Philosophy of Science* for extraordinarily generous and helpful comments.

Department of Philosophy
University of North Carolina Chapel Hill
NC 27599-3125, USA
ramneta@yahoo.com

References

- Austin, J. L. [1964]: *Sense and Sensibilia*, New York: Oxford University Press.
- Boghossian, P. [2006]: *Fear of Knowledge*, Oxford: Oxford University Press.
- Christensen, D. [1997]: 'What is Relative Confirmation?', *Noûs*, **31**, pp. 370–84.
- Davidson, D. [2001]: 'A Coherence Theory of Truth and Knowledge', in *Subjective, Intersubjective, Objective*, Oxford: Clarendon Press.
- Fantl, J. and McGrath, M. [2002]: 'Evidence, Pragmatics, and Justification', *Philosophical Review*, **111**, pp. 67–94.
- Feldman, R. and Conee, E. [1985]: 'Evidentialism', *Philosophical Studies*, **48**, pp. 15–34.
- Feldman, R. [1988]: 'Having Evidence', in D. F. Austin (ed.), *Philosophical Analysis*, Dordrecht: Kluwer.

- Joyce, J. [2004]: 'Williamson on Knowledge and Evidence', *Philosophical Books*, **45**, pp. 296–305.
- Kaplan, M. [1996]: *Decision Theory as Philosophy*, Cambridge: Cambridge University Press.
- Lewis, D. [1996]: 'Elusive Knowledge', *Australasian Journal of Philosophy*, **74**, pp. 549–67.
- Moore, G. E. [1993]: 'Proof of an External World', in T. Baldwin (ed.), *G.E. Moore: Selected Writings*, London and New York: Routledge.
- Neta, R. [2007]: 'Propositional Justification, Evidence, and the Cost of Error', *Philosophical Issues*, **17**, pp. 197–216.
- Plantinga, A. [1993]: *Warrant: The Current Debate*, Oxford: Oxford University Press.
- Rorty, R. [1979]: *Philosophy and the Mirror of Nature*, Princeton, NJ: Princeton University Press.
- Schaffer, J. [2005]: 'Contrastive Knowledge', in T. Gendler and J. Hawthorne (eds), *Oxford Studies in Epistemology, volume 1*, Oxford: Oxford University Press.
- Stanley, J. [2005]: *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Williamson, T. [2000]: *Knowledge and its Limits*, Oxford: Oxford University Press.