

Defeating the Dogma of Defeasibility

Ram Neta

Ever since Gettier 1963 convinced English-speaking philosophers that justified true belief does not suffice for knowledge, many epistemologists have been searching for the elusive “fourth condition” of knowledge: the condition that must be added to justification, truth, and belief, in order to get a set of non-trivial conditions that are individually necessary and jointly sufficient for knowledge.¹ The problem of finding such conditions is generally known as the “Gettier problem”. Many different fourth conditions have been proposed and subsequently counterexamples, and some philosophers have suggested that some of the other three conditions may need to be revised as well. But recently, several philosophers have suggested that the Gettier problem is insoluble, and that the theory of knowledge should proceed in a more modest way: not by trying to specify a set of non-trivial conditions that are individually necessary and jointly sufficient for knowledge, but rather by trying to figure out whatever we can about knowledge, and in particular to figure out the role that knowledge, and knowledge ascription, play in our lives. Williamson 2000 has developed an epistemological view that falls within this latter, more modest approach to the theory of knowledge, and his work has been immediately and justly influential.

Now, it is all to the good to figure out whatever we can about knowledge, and to examine the role that knowledge, and knowledge ascription, play in our lives. But we should also try to understand what it is about knowledge and knowledge ascription that suit them to play those roles. Suppose that, as Williamson says, it is constitutive of the

act of assertion that one is entitled to assert only what one knows to be true. Or suppose that, as Hawthorne 2003 says, one can employ a premise in practical reasoning only if one knows the premise to be true. We might wonder what it is about knowledge that renders it suitable to play this constitutive role in assertion or in practical reasoning.

What is it about knowing that makes it the case that one is one entitled to assert only what one knows to be true, or one is entitled to reason practically from nothing other than what one knows to be true? What's so special about knowing that enables it to play these roles? These seem to be perfectly good questions, to which we should like answers.

Of course, if the nature of knowledge were exhausted by its role in assertion, or in practical reasoning, then the questions above could be answered trivially: what's special about knowing would be precisely that it has this role in assertion, or in practical reasoning, and there would be no further fact about knowing by virtue of which it has these roles – there would be nothing that *renders* knowledge suitable to play these roles, since knowledge would be *simply* whatever plays these roles. Now, in fact, this is not true of knowledge. Knowledge may be, as Williamson thinks, the constitutive norm of assertion, or it may be, as Hawthorne thinks, the norm for premises of practical reasoning. But that is not, on anyone's view, *all* that knowledge is. There is more to say about what knowledge is. And, it is fair to hope, some of what more we can say about what knowledge is can help us to understand what renders it suitable to play the role(s) it plays in our lives – for instance, if Williamson is right, to be the norm of assertion.

One way – of course not the only way – that we might try to answer questions about what renders knowledge suitable to play the role that it plays in our lives, and what makes it the case that knowledge has the various properties that it has, is by giving an

informative account of *what knowledge is*, and then appealing to that account in order to explain various features of knowledge. Just as we can explain various features of water by appeal to an informative account of what water is, and we can explain various features of zebras by appeal to an informative account of what zebras are, and we can explain various features of inflation by appeal to an informative account of what inflation is, so too – one might hope – we might be able to explain various features of knowledge by appeal to an account of what knowledge is.

In order to offer a correct and informative account of what knowledge is, I claim that we will need to give up one of Williamson's main contentions – namely, that for any epistemic agent *S*, *S*'s evidence set (at a time *t*) comprises all and only those propositions that *S* knows (at *t*) to be true. If we give up this equation of evidence and knowledge, we will then be in a position to accept a particular account of knowledge as both true and informative – and, as we shall see, the account that I offer below helps to substantiate some other distinctive Williamsonian epistemological theses (e.g., that if *S* knows that *p*, then the conditional probability of *p* on *S*'s evidence set is 1; and that the preceding point does not entail that *S* be rationally required to be completely confident that *p* is true). If we continue to accept Williamson's equation of evidence and knowledge, then the account of knowledge that I offer below may not seem to be informative, since that account explains what it is for *S* to know that *p* partly in terms of *S*'s evidence.²

I accept a particular account of what knowledge is, and I propose to state some of its main features here. The account of knowledge that I accept implies that knowledge is indefeasible, and so it will strike virtually all epistemologists as suffering from obvious and devastating problems. For instance, Williamson offers an example to illustrate the

defeasibility of knowledge, and his example will initially strike many philosophers as a clear counterexample to my account of knowledge. I devote section III of this paper to showing why this and other seemingly obvious counterexamples to the IJTB account of knowledge (specifically, to its indefeasibility) do not tell against that account.

But before proceeding to defend my account (which I call the IJTB account) against these various objections, I will devote section I to laying out some competitor accounts, and presenting counterexamples to each of those accounts. Of course, my survey of competitor accounts in section I will not aspire to historical completeness: such an aspiration could be realized only by a very lengthy book. Rather than identify all, or even most, of the actual competitors to the IJTB theory of knowledge, I will simply describe some simple competitor accounts that can be formed by adding to the JTB account various possible “fourth conditions”, each of which has made its way into one or another prominent theory of knowledge. I have two reasons for doing this. First, I want to show that it is not at all unreasonable to despair of giving a correct and informative account of knowledge: the history of such attempts is not promising. Although I think we should not give up, it is easy to understand how someone might reasonably think otherwise. Second, I want to motivate the IJTB account by showing that it delivers the right verdicts about various cases described in section I, and that it identifies the kernel of truth in each of the competitor accounts described in section I.

In short, I hope to show that the traditional immodest program in the theory of knowledge can be brought to fruition, despite the bleak history of attempts to do so. We need not confine ourselves to finding out various features of knowledge and of knowledge ascription, but we can explain those various features of knowledge and of

knowledge ascription by appeal to a true and informative answer to the question “what is knowledge?”

What we want from an account of knowledge, I take it, is this: The account should predict that, and explain why, clear cases of knowledge are indeed cases of knowledge; the account should predict that, and explain why, clear cases of non-knowledge are indeed cases of non-knowledge; the account should predict that, and explain why, unclear or disputable cases of knowledge are indeed unclear or disputable; and finally, the account should help us to understand how it is that knowledge and knowledge ascription can play the role(s) that they play in our lives. I see no reason to demand that an account reduce knowledge to non-epistemic, let alone non-normative, phenomena. It is good enough if we can explain knowledge in terms of things other than knowledge, so long as we do not have to invent new primitives, or appeal to any notions that philosophers do not *already* need.

For the purposes of this paper, I will start by assuming that a true account of knowledge is of the form:

S knows that p = S has a justified, true belief that p and ...

where the dots are to be filled in by a specification of the fourth condition of knowledge, the anti-Gettier condition. It will eventually turn out that the satisfaction of the fourth condition will entail the satisfaction of two of the three standard conditions (justification and truth). Of course, there are plenty of questions that arise in connection with the three specified conditions on knowledge – i.e., justification, truth, and belief. For instance, we

might wonder what sorts of things can serve to confer justification (by which I here mean *doxastic justification*, not merely propositional justification) upon a belief (Other beliefs? Experiences? Being caused or sustained by a reliable process? Being sufficiently well supported by evidence given one's present purposes?) and we might wonder what sort of structure there is among such justifiers (foundationalist? coherentist? infinitist?). Again, we might wonder what relation there must be between a belief and the thing that confers justification upon it in order for the belief to be justified (i.e., the nature of the "basing relation"). And of course, there are any number of venerable questions about the nature of truth and the nature of belief. I will not address any of these important questions in this paper. I confine myself to defending an account of knowledge that fills in the fourth condition above, and I will not say anything more in this paper about the three standardly accepted conditions.

I. What is Knowledge? (In Search of the Elusive Fourth Condition)

Before Gettier's paper, one or another version of the following theory of knowledge was accepted by C.I. Lewis, A.J. Ayer, Roderick Chisholm, and others³:

JTB Theory:

S knows that p = S has a true, justified belief that p.

But there are counterexamples to the JTB theory, such as the following:

Smith's evidence clearly indicates that his coworker Nogot owns a Ford, and so Smith justifiably believes that Nogot owns a Ford. From this belief Smith infers that one

of his coworkers owns a Ford. Unbeknownst to Smith however, Nogot does not own a Ford, but another one of Smith's coworkers, Havit, does own a Ford. So Smith has a justified true belief that one of his coworkers owns a Ford, but he doesn't know that one of his coworkers owns a Ford.

The JTB theory is false. Since this was discovered⁴, many different accounts of knowledge have been proposed. In this section, I briefly review a small sampling of those proposed accounts, and then give counterexamples to each of them. I make use of these counterexamples in the next section, when I show that my own IJTB account of knowledge delivers the right verdicts about all of these examples.

Here, then, are some accounts of knowledge that have been offered in the wake of the demonstrated failure of the JTB theory:

No-False-Lemmas Theory:

S knows that $p = S$ has a justified true belief that p and S 's belief that p is not based upon reasoning to p from any false premises.⁵

Counterexample: Latifa knows, on the basis of countless pieces of journalistic evidence, that U.S. troops are in Iraq. But she bases her belief that U.S. troops are in Iraq on many different inferences that she makes, from many different premises. Most of those inferences are sound, and the premises are known by her to be true, but not all of the inferences are like that. For instance, she infers that U.S. troops are in Iraq from the false proposition that U.S. troops have discovered weapons of mass destruction in Iraq.

Her knowing that U.S. troops are in Iraq is thus compatible with her basing her belief upon – *inter alia* – reasoning from a false premise.

Causal Theory:

S knows that p = S has a justified true belief that p, and S is caused to believe that p (in the way that S believes it) by the fact that p.⁶

Counterexample: Henry is driving through the countryside, sees a barn, and thereby comes to believe that he sees a barn. He has a justified true belief that he sees a barn, and he is caused to believe that he sees a barn by the fact that he sees a barn. Nonetheless, he doesn't know that he sees a barn, since most of the apparent barns in this area are not barns at all, but only barn facades.

Explanation Theory:

S knows that p = S has a justified true belief that p and the fact that p explains S's believing that p (in the way that S believes it).⁷

Counterexample: Rachel owns a Ford, but she doesn't want anyone to know that she owns a Ford. So she produces lots of evidence that the Ford belongs to Isaac. Rebecca thus forms a justified belief that Isaac owns a Ford, and infers from this the true conclusion that someone in her office owns a Ford. Rebecca thus comes to have a justified true belief that someone in her office owns a Ford, and the fact that someone in her office owns a Ford explains Rebecca's belief that someone in her office owns a Ford. Nonetheless, Rebecca doesn't know that someone in her office owns a Ford.

Sensitivity Theory:

S knows that p = S has a justified true belief that p and S wouldn't believe that p in the way that S does if p were not true.⁸

Counterexample: Bridget is driving through the countryside, sees a cow, and thereby comes to believe that that particular object (pointing to the object in question) is a cow. She thereby has a justified true belief that that particular object is a cow. But there is no possible world in which that same particular object (or its counterpart, if it has one) is not a cow. Thus, at least according to the standard Lewis-Stalnaker semantics for counterfactuals⁹, it is vacuously true that she wouldn't believe that that particular object is a cow if it were not a cow. Nonetheless, Bridget doesn't know that that particular object is a cow, since most of the apparent cows in this area are not cows at all, but rather cow facades.

One might worry that the counterexample depends upon a particular, contentious account of the semantics for counterfactuals. But no matter what account we give of the semantics for counterfactuals, it is at least not clearly false that Bridget wouldn't believe that that particular object is a cow if it were not a cow. But, according to the Sensitivity Theory, this would have to be clearly false in order for the case of Bridget to be (as it is) clearly such that Bridget doesn't know that that particular object is a cow. So, even if the Lewis-Stalnaker semantics for counterfactuals is wrong, the case still provides a counterexample to the Sensitivity Theory.

Safety Theory:

S knows that $p = S$ has a justified true belief that p and there is no nearby possible world in which S , forming her belief in the way that she does, falsely believes that p .¹⁰

Counterexample: In the course of being taught some rudimentary anatomy, Jacques seems to hear his teacher tell him that human beings have brains, and Jacques thereby comes to have a justified true belief that he, Jacques, has a brain. Furthermore, this belief is safe, because there is no possible world (a fortiori no nearby possible world) in which Jacques holds the belief but it is false. Nonetheless, Jacques fails to know that he has a brain, because his auditory experience as of his teacher telling him that he has a brain is illusory: in fact, his teacher was deceptively telling Jacques that human beings do not have brains, but Jacques simply misheard what the teacher said.

Indefeasibly Justified True Belief Theory:

S knows that $p = S$ has a justified true belief that p and there is no true proposition q which is such that, if S were to be justified in believing that q , then S would no longer be justified in believing that p .¹¹

Counterexample: Siobahn believes that Elvis is still alive, and she knows that she believes that Elvis is alive. But, if Siobahn were to be justified in believing that Elvis is dead, then she would not believe that Elvis is alive, and so she would then not be justified in believing that she believes that Elvis is alive. The counterexample here is induced by the use of the counterfactual conditional in the theory.¹²

No-Justified-Falsehoods Theory:

S knows that $p = S$ has a justified true belief that p and

there is no true proposition e such that the total evidence set E' (where $E' =$ the conjunction of e and S 's actual evidence set E) does not constitute a justification for S to believe that p , and also does not constitute a justification for S to hold any false beliefs.¹³

Counterexample: I know lots of empirical facts, e.g., that George W. Bush is the President of the United States. But take any proposition which is both true and extremely unlikely to be true, for instance the proposition Z that states precisely the actual distribution of microphysical properties over points in space-time. Z is true, and therefore the following disjunction is true as well: George W. Bush is not the President of the United States, or Z is true. But my total evidence constitutes a justification for me to believe that this disjunction is false, and so the conjunction of my total evidence with the true proposition that, say, I am now thinking, constitutes a justification for me to believe a falsehood.

In general, whenever S has total evidence E , and S knows that p , there will be some extremely improbable truth which is such that S is justified in believing that it is false. (If we think that evidential justification is probabilistic, then we can put the point this way: there will in general be some true proposition q which is such that $\text{Prob}(q/E)$ will be extremely low, and perhaps even lower than $\text{Prob}(q)$.) And if there is such a true proposition q , then E constitutes a justification for S to believe the falsehood that not- q . If that is correct then, no matter what the value of E , it will constitute a justification for S to believe some falsehood.

Having briefly disposed of the accounts of knowledge given above, in the next section I will offer a first approximation to my own account of knowledge – the

Indefeasible Justified True Belief account – and then show how it can handle cases that serve as counterexamples to other accounts.

II. Knowledge is Indefeasible Justified True Belief

Here is a first approximation to the account of knowledge that I will defend in this paper – what I call the IJTB account of knowledge. (I will provide a fully accurate statement of the account not just yet, but only in section IV. To start with, I'll treat the following statement as accurate, and refine it only later.)

S knows that p = S has a justified, true belief that p , and there is no true proposition e such that the conjunction of e and S's actual evidence set E does not constitute a justification for S to believe that p .

Proposals similar to this one were considered early in the post-Gettier literature, but they were immediately dismissed on the allegedly obvious grounds that they are too strong. Knowledge, it is commonly thought, is at least sometimes defeasible by future evidence, and the IJTB account proposed above is incompatible with this obvious fact. Thus – it is inferred – the IJTB account is too strong. In the next section, I will critically examine various putative counterexamples to the IJTB account of knowledge. Before proceeding to answer those alleged counterexamples to IJTB, I will devote the present section to showing how IJTB is strong enough to handle all of the examples of non-knowledge that end up counting as knowledge according to the theories listed in the previous section.

The case of Havit's Ford: Though Smith has a justified true belief that someone in his office owns a Ford, the conjunction of his current evidence with the true proposition that Nogot does not own a Ford does not constitute a justification for Smith to believe that someone in his office owns a Ford. (A perfectly analogous treatment applies to the case of Rachel's Ford.)

The case of the barn facades: Though Henry has a justified true belief that he sees a barn, the conjunction of his current evidence with the true proposition that most of the apparent barns in this area are not barns does not constitute a justification for Henry to believe that he sees a barn. (A perfectly analogous treatment applies to the case of Bridget and the cow facades – and this is as it should be, given that the cases seem to be epistemically alike, even if metaphysically different.)

The case of Jacques's brain: Though Jacques has a justified true belief that he has a brain, the conjunction of his current evidence with the true proposition that he misheard his teacher's lecture does not constitute a justification for Jacques to believe that he has a brain.

The case of Siobahn's second-order belief: If Siobahn were to be justified in believing that Elvis is dead, then she would not believe that Elvis is alive, and so would not believe that she believes that Elvis is alive. But the conjunction of Siobahn's evidence set (including the fact that she believes that Elvis is alive) with the true proposition that Elvis

is dead constitutes a justification for Siobahn to believe that she believes that Elvis is alive.

The IJTB account smoothly handles all of these examples. Furthermore, it specifies the kernel of truth in each of these other accounts. If it's true that you would believe that p even if p were false, then the conjunction of that truth with your total evidence will typically fail to justify you in believing that p. If it's true that your belief that p is based on your inferring p from a false premise, then the conjunction of that truth with your total evidence will typically fail to justify you in believing that p. If it's true that your belief that p is not caused by, or explained by, the fact that p, then the conjunction of these truths with your total evidence will typically fail to justify you in believing that p. And so on.

So the IJTB has something to be said for it. But isn't the IJTB account subject to counterexamples of its own? In the next section, after considering some such examples, I argue for a negative answer to this question.

III. Apparent counterexamples to the IJTB theory of knowledge

Many philosophers would object that there are obvious and devastating counterexamples to the IJTB theory. In this section, I critically examine those alleged counterexamples.

Williamson offers the following example to illustrate what he calls the "defeasibility" of knowledge:

“I see one red and one black ball put into an otherwise empty bag [call this ‘e’] Now suppose that on the first ten thousand draws a red ball is drawn each time, a contingency which my evidence does not rule out in advance, since its evidential probability is non-zero. But when I have seen it happen, I will rationally come to doubt e; I will falsely suspect that the ball only looked black by a trick of the light.”¹⁴

Now, what does this example show? It shows that it’s possible for the following sort of thing to occur: at a particular time, S knows that p on the basis of evidence e, and subsequently S gains additional evidence e’ and thereby loses her knowledge that p. Knowledge is clearly, in this sense, defeasible.

But notice that S’s gaining evidence e’ might result in S’s losing her belief that p is true, or it might result in S’s losing some of her original evidence for p, or it might result in S’s losing the ability *reasonably to form* her belief that p on the basis of her evidence (even if her total evidence continues to constitute a justification for her to believe that p). On any of these last three scenarios, S’s gaining evidence e’ will result in S’s losing her knowledge that p. But all this is compatible with IJTB. IJTB says nothing about what *would* happen to our epistemic subject if she were to gain an additional bit of evidence. It says that, if S knows that p, then, for any true proposition e’, the conjunction of our subject’s actual evidence set with e’ constitutes a justification for our subject to believe that p.

Now, Williamson himself might agree with the latter claim, since he thinks that $E=K$, and so whenever S knows that p, p itself is included in S’s evidence set. But if p is in S’s evidence set, then for any true proposition q, the conjunction of q with S’s

evidence set will constitute a justification for S to believe that p, since p is implied by that conjunction. So, by Williamson's own lights, his example of the red ball would not serve as a counterexample to IJTB. Indeed, since Williamson accepts the claim that your evidence set includes all and only what you know to be the case, he might accept that the IJTB account (as presented so far) is true, but he might not regard it as informative.

But I reject Williamson's thesis that $E=K$. Indeed, I reject even the weaker thesis that your evidence set includes all and only what you know to be the case. The latter thesis is both too strong a condition on evidence, and also too weak a condition on evidence. It's too strong because sometimes p is in my evidence set even though I am not aware that p. For instance, at present I am visualizing a speckled hen, and my evidence set includes the fact that *the speckled hen image contains 7 speckles*. (Seven is just a big enough number that I might *fail* to notice that there are that many speckles, but it is just a small enough number that I am able simply to *notice* that there are that many speckles – I don't need to employ any arithmetical procedures in arriving at the belief that there are that many speckles.) Whether I know it or not, I have evidence that gives me conclusive justification for believing that the image contains 7 speckles: namely, the fact that the image does indeed contain 7 speckles. That fact is in my evidence; it makes a difference to how I rationally ought to distribute my confidence over hypotheses. Some of the truths that are in my evidence set may be unknown by me, even if each one of those truths is knowable by me upon the right kind of reflection. And of course, it's possible that each one of those truths is reflectively knowable by me even if all of them aren't simultaneously knowable by me.

Williamson's account is too weak a condition on evidence because it allows that everything I know by deduction from my evidence is in my evidence set. But this is not true. Here's why: for any proposition *q* that is in my evidence set, if you ask me what evidence I have for *q*, I am entitled to reject your request for evidence, or else answer you simply by citing *q* itself. In other words, for it to be rational for me to believe *q* – if *q* is in my evidence set – I do not need to base my belief that *q* upon any evidence other than *q*. Now suppose that I know that *p*, and I know it *solely* because I've competently deduced *p* from premises *a*, *b*, and *c* (each of which I know to be true), and then you ask me what evidence I have for *p*. In such a case, I typically cannot give you a true and complete answer without citing *a*, *b*, and *c*. But if *p* itself were in my evidence set, then I could give you a true and complete answer to the question what evidence I have for *p* simply by citing *p*. Since there are at least some occasions on which I cannot do this, *p* is not (at least on those occasions not) in my evidence set, though I know it to be true. Thus, not everything that I know by deduction from my evidence is itself in my evidence set.

My evidence set at a particular time includes all and only those truths that I am entitled to take for granted in inquiry at that time.¹⁵ Now, what is it for *S* to be entitled to take it for granted in inquiry that *p*? To be entitled to take it for granted that *p*, in the sense in which I'm using that phrase here, involves at least this much: being entitled to appeal to *p* *without having to be able to defend p against challenges*. There are lots of things that I am now entitled to appeal to in the course of inquiring whether or not, say, it will rain tomorrow. But not all of these things that I'm entitled to appeal to are things that I'm entitled to *take for granted*; some of the things that I'm entitled to appeal to are

things my entitlement to appeal to which requires that I be able to defend them against reasonable challenges. For instance, in predicting whether or not it will rain tomorrow, I might be entitled to appeal to the claim that the warm front that's moving towards us from the west will probably reach here by tomorrow morning – but I would be entitled to appeal to this claim only if I could defend it against reasonable challenges, such as “how do you know that it won't slow down and not reach here until tomorrow night?” So, while I'm entitled to appeal to the claim about the warm front in inquiring whether or not it will rain tomorrow, I'm not entitled to *take it for granted*, in the relevant sense. So it's not in my evidence set, as I conceive of it. My evidence set will include only those truths that I'm entitled to appeal to even if I cannot defend them against challenges. That might include statements like “I am now being appeared to redly”, but it might also include statements about the physical world like “we are on the planet Earth”. These are truths that I'm entitled to appeal to even if I cannot defend them against challenges. If someone challenges my claim that I am now being appeared to redly, or my claim that we are on the planet Earth, then I would probably have no idea how to defend my claims against those challenges, nor would I need to do so. I would simply dismiss the challenges as silly. Among the things that I'm entitled to appeal to in judgment, the only ones that count as in my evidence set are those things that I am entitled to appeal to *whether or not* I am able to defend them against challenges.¹⁶

This conception of evidence may seem more restrictive than our ordinary conception of evidence. If the weatherman tells me that there is a 70% chance of rain tomorrow, and then I ask him for his evidence, it may seem perfectly reasonable and correct for him to cite as his evidence the following fact: the warm front that's moving

towards us from the west will probably reach here by tomorrow morning. But if I challenge his belief that the warm front that's moving towards us from the west will probably reach here by tomorrow morning, doesn't he need to be able to defend this belief against this challenge? He can't simply dismiss my challenge as silly. Doesn't he need to be able to justify his belief, in order to reasonably appeal to it? If the answer to this question is "yes", then I do not count it as in his evidence that the warm front that's moving towards us from the west will probably reach here by tomorrow morning. What I count as his evidence in that case will then include the things to which he appeals in justifying his belief that the warm front that's moving towards us from the west will probably reach here by tomorrow morning – at least if those things are such that he can appeal to them without having to defend them against challenges. (These things might include the testimony of the meteorologist who feeds information to the weatherman, or they might include the readings of various measuring instruments that the weatherman consults, or they might include the weatherman's perceptual experiences.)

This suggests two tests that we might usefully employ in figuring out whether S is, at a certain time, entitled to take it for granted in inquiry that e. First: if S can, at time t, truthfully and relevantly answer a question of the form "how do you know that e?" by saying simply "e" – or alternatively, if S is entitled to dismiss the question "how do you know that e?" – then this is one indication that S is entitled, at t, to take e for granted in inquiry. Second, if S can, at t, truthfully and relevantly answer other questions of the form "how do you know that p?" by saying "e, and e supports p", then this is a second indication that S is entitled, at t, to take e for granted in inquiry. These are all the general

remarks that I will offer here about what it is for a proposition to be in one's evidence set.¹⁷

Let's consider how this account of evidence applies to cases. Consider the case of Henry driving through barn façade country. What is Henry entitled to take for granted in inquiry? That depends on the details of the case, but the following would seem to be correct on any natural way of filling in the case. Henry is entitled to take it for granted that he sees something that looks like a barn. If asked how he knows that he sees something that looks like a barn, Henry could truthfully and relevantly answer by saying something like: "What a silly question! I just *do* see something that looks like a barn". And if asked how he knows, say, that there are any objects that look like barns in this region, Henry could truthfully and relevantly answer by saying that he sees something that looks like a barn. But Henry is not entitled to take it for granted that he sees a barn. If asked how he knows that he sees a barn, Henry could not truthfully and relevantly answer by saying something like "what a silly question! I just do see a barn". It is, of course, *true* that Henry sees a barn, but, in the case envisaged, this is not a relevant answer to the question.

All of the above seems to me correct about the case of Henry driving through barn façade country. But we can generate different verdicts about what Henry is or is not entitled to take for granted in inquiry by ringing changes on features of the case. For instance, if Henry is driving through a countryside in which there are hallucination-inducing drugs floating in the air, then Henry might not be entitled to take it for granted that he sees something that looks like a barn. In such a case, he might be entitled to take it for granted merely that *it looks to him as if* there is a barn before him. But in a case in

which there are no such epistemic risks lurking – no barn facades, no hallucination-inducing drugs, nothing of the sort – Henry might be entitled to take it for granted that he sees a barn. I hope that these remarks about what someone is or is not entitled to take for granted in inquiry suffice for present purposes.

Given his own view that $E=K$, Williamson himself would not regard his own example of the black ball as a counterexample to the IJTB theory, but many other philosophers (those of us who reject Williamson's thesis that $E=K$) are likely to regard Williamson's example as a counterexample to the IJTB theory. Is it indeed a counterexample to IJTB? Before the ten thousand draws are made, but once I've seen the black ball in the bag, I know that

(p) there is a black ball inside the bag,

and I know that p on the basis of some visual evidence V – never mind precisely how to specify the propositional content of V (perhaps it is *that there is a black ball in the bag*, or it is *that I see a black ball in the bag*, or it is *that it looks to me as if there is a black ball in the bag*, etc.). Now after the ten thousand drawings are made, I have gained some additional evidence concerning the track-record produced by the series of draws – call this track-record evidence T. Now, would I be justified in believing that p on the basis of the conjunction of V and T? Well, imagine that at a particular moment, I have both pieces of evidence in my evidence set: I have visual evidence V *and* I have the track-record evidence T. (We can imagine what it would be like to have this conjunction in our evidence set all at once by imagining that the bag is transparent.) This conjunction of V

and T justifies me in believing that p. In that case, why is it that, in Williamson's example, once I have made the ten thousand drawings, then I am no longer justified in believing that p (i.e., there is black ball in the bag)? It is because my memory of V is not as good evidence for p as is V itself. Indeed, it normally tends to happen that, as the time at which I have evidence V recedes farther and farther into the past, the strength of my evidence for p grows weaker and weaker. Williamson's example is not a counterexample to the IJTB theory.¹⁸

Harman 1968 and Klein 1971 each describe other cases that might be thought to be counterexamples to the IJTB theory. I will now describe these cases along with a further apparent counterexample, and then describe a recipe for generating further such apparent counterexamples to IJTB. Finally, I will argue, on grounds that are independent of IJTB, that these cases cannot serve as counterexamples to IJTB.

The case of the retracted newspaper story: A famous official has just been assassinated, and you read about the assassination on the front page of a major newspaper. You thereby acquire knowledge that the official has been assassinated. The next day, the government decides to cover up the assassination, and forces all the newspapers to run a retraction of the earlier story, claiming that the official is doing fine, and the assassination attempt was not successful. You hear nothing about the retractions, and continue to have a true belief that the official was assassinated, but your knowledge has (allegedly) been defeated. (Harman 1968, 172, with slight modifications)

The case of the library thief and his lying mother: You've just seen Tom Grabit steal a book from the library. You know Tom Grabit reasonably well, and are therefore justified in believing that it is Grabit whom you saw. You now know that Grabit stole a book from the library. Once she finds out about the episode, though, Tom's mother tells the authorities that Tom has an identical twin brother John, and it was John who stole the book from the library. In fact, Tom has no brother, and his mother made up the whole story in order to protect Tom from prosecution. But your knowledge has (allegedly) been defeated. (Klein 1971, 474, with several modifications)

The case of the A's that are mostly B's: Suppose that almost all of the 1,000,000 A's are B's, and I have seen one hundred A's, and all of those A's have been B's. It may seem that this suffices, at least under certain conditions, for me to know that most A's are B's. But what if, among the 1,000,000 A's, there are 1,000 A's that are not B's. Does the conjunction of the true proposition that there are 1,000 A's that are not B's, along with my current evidence, still justify me in believing that most A's are B's? It may seem clear that it does not, for I have observed only 100 A's so far – and so this case seems to be a counterexample to the IJT theory.

All of the three examples above share a common profile, and we can appeal to this profile in order to construct a general recipe for generating putative counterexamples to IJT, a recipe adapted from Harman 1973. Suppose that S knows that p, and knows it on the basis of her evidence e, which is such that $\text{Prob}(p/e) < 1$. Let Z be a true proposition which is very unlikely to be true, given S's total evidence e: for instance, let

Z be a true proposition specifying precisely the distribution of microphysical states over all points in space-time. The disjunction (Z or not-p) is true, since the first disjunct is true. Now does the conjunction of that disjunction (Z or not-p) with e constitute a justification for S to believe that p? It may seem that it cannot do so, since Prob (p/e & (Z or not-p)) is very low. (In a chancy world in which some quantities vary continuously, we can pick a true Z which is as unlikely as you please.) But then the IJTB theory rules against the possibility of knowing p on the basis of evidence e if Prob (p/e) < 1. Is that not a reductio of the IJTB theory?

We can state the argument above more precisely: if we let q = the disjunction (Z or not-p), and if we let N be an arbitrarily large finite number such that Prob (Z) < 1/N, and if we suppose that, for S, Prob (p) > 0, Prob (p/e) < 1, and 0 < Prob (q) < [Prob (p) x Prob (e & not-p)], then:

- (1) Prob(q) < Prob(p)Prob(e¬-p) (follows from the stipulations)
- (2) Prob(p/e&q) = Prob(p&e&q)/Prob(e&q) (definition of Prob)
- (3) (p&e&q) is logically equivalent to (p&e&Z).
- (4) Prob(p/e&q) = Prob(p&e&Z)/Prob(e&q). (from 2, 3)
- (5) (p&e&Z) logically entails Z.
- (6) Prob(p&e&Z) < or = Prob(Z) = 1/N (from 5)
- (7) (e¬-p) logically entails (e&q). (from definition of q)
- (8) Prob(e¬-p) < or = Prob(e&q) (from 7)
- (9) Prob(p/e&q) < or = (1/N)/Prob(e&q) (from 4, 6)
- (10) (1/N)/Prob(e&q) < or = (1/N)/Prob(e¬-p) (from 8)

(11) $(1/N)/\text{Prob}(e \& \text{not-}p) < \text{Prob}(p)\text{Prob}(e \& \text{not-}p)/\text{Prob}(e \& \text{not-}p)$ (from 1)

(12) $\text{Prob}(p)\text{Prob}(e \& \text{not-}p)/\text{Prob}(e \& \text{not-}p) = \text{Prob}(p)$

(13) $\text{Prob}(p/e \& q) < \text{Prob}(p)$ (from 9-12)

So what we can conclude from this is that, if S knows that p on the basis of evidence e, where $\text{Prob}(p/e) < 1$, and then – given that there are many truths that are very unlikely to be true, given our total evidence – there will be a true proposition Z such that $\text{Prob}(p/e \& Z)$ is very low, and much lower than $\text{Prob}(p)$. And so there will be a true proposition Z such that e&Z does not constitute a justification to believe that p. This seems decisively to refute the IJT theory of knowledge, at least if S can know that p on the basis of evidence e such that $\text{Prob}(p/e) < 1$.

But this is the crucial question: is it possible for S to know that p on the basis of evidence e if $\text{Prob}(p/e) < 1$? In the next section, I will argue that it is not.

IV. Knowledge Requires Belief Based on Infallible Evidence

In this section, I will argue that, if S knows that p on the basis of evidence e, then, for S, $\text{Prob}(p/e) = 1$. Indeed, I will argue for an even stronger conclusion than this, namely, that S cannot know that p on the basis of evidence e unless there is no *possible* situation in which S has e and yet p is not true.

My argument will employ the following premises:

Two-premise Closure: If S knows that p, and S knows that q, and S competently deduces r from p and q, while retaining her knowledge that p and her knowledge that q, then S knows that r.

Transmission of Empirical Evidential Basis: If S knows that p solely on the basis of S's empirical evidence e, and if S knows that q solely by deducing q from the conjunction of p and e, then S knows that q on the basis of empirical evidence e.

Now, I will argue that, if S knows that e, and S knows that p on the basis of e, then e is infallible evidence for p, i.e., e is such that S cannot have e unless p is true.

Consider the following facts:

I cannot know that the paper I'm reading is not in error, solely on the basis of the evidence that the paper says that a famous official was assassinated.

I cannot know that the 100 A's I've observed are representative of A's with respect to their B-hood solely on the basis of the evidence that A1 is a B, A2 is a B, ..., A100 is a B.

I cannot know that my visual experience as of Tom stealing the library book is accurate solely on the basis of the visual experience as of Tom stealing the library book.

There are, as a matter of empirically obvious fact, countless more facts of the kind above, in which *e* is fallible evidence for some proposition *p*, but *S* cannot know, on the basis of *e* alone, that *e* is not misleading with respect to *p*. A survey of such cases inductively justifies the following conclusion: if *e* is fallible evidence for *p*, then one cannot know, on the basis of *e* alone, that *e* is not misleading with respect to *p*.

Now consider the following additional facts:

I cannot know *a priori* that the paper that I'm reading is not in error in saying that the famous official was assassinated.

I cannot know *a priori* that the 100 A's that I've observed are typical of A's with respect to their B-hood.

I cannot know *a priori* that my visual experience of Tom stealing the book is veridical.

Again, there are, as a matter of empirically obvious fact, countless more facts of the kind above, in which *e* is fallible evidence for some proposition *p*, and *S* cannot know *a priori* that *e* is not misleading with respect to *p*. A survey of such cases inductively justifies the following conclusion: if *e* is fallible evidence for *p*, then one cannot know *a priori* that *e* is not misleading with respect to *p*.

Now, by two-premise closure, if *S* knows that *p* on the basis of empirical evidence *e*, and if *S* deduces from *e* and *p* that *e* is not misleading with respect to *p*, then *S* has knowledge that *e* is not misleading with respect to *p*. If *S* has this knowledge *by*

deducing it from e and p , then, by the transmission of empirical evidential basis, S has this knowledge on the basis of evidence e alone. But our first survey of cases above showed that, if e is fallible evidence for p , then S cannot know that e is not misleading with respect to p solely on the basis of e . And our second survey of cases above showed that, if e is fallible evidence for p , then S cannot know a priori that e is not misleading with respect to p . So, if S knows that e is not misleading with respect to p , then either S knows this on the basis of some evidence e' that is distinct from e , or else e is not fallible evidence for p .

Let's consider the first disjunct. If S knows, on the basis of e' , that e is not misleading with respect to p , then how might S know that the conjunction e and e' is not misleading with respect to p ? Does S know this on the basis of some further evidence e'' ? The appeal to further evidence e' simply sets us on a regress. I conclude that the second disjunct is true: if S knows that e is not misleading with respect to p , then e is not fallible evidence for p . And so, if S knows that p on the basis of evidence e , then e is not fallible evidence for p . For e to give me knowledge that p , e must be the kind of evidence which is such that, on the basis of that evidence alone, I can know that it is not misleading with respect to p . The survey of cases above inductively confirms the hypothesis that no fallible evidence satisfies this latter description. We have thus shown, by means of this inductive argument, that fallible evidence cannot give one knowledge.¹⁹

This last argument provides all the resources necessary to handle proposed counterexamples to the IJT theory. If it's a necessary condition of S 's knowing that p that S 's evidence e for p be infallible – i.e., such that there is no possible situation in which S has e but p is not true – then the conjunction of e with anything constitutes a

justification for S to believe that p. This guarantees that we will not be able to find cases of knowledge that fail to satisfy the controversial fourth condition of the IJTB theory.

Of course, if Williamson is right that $E=K$, then of course knowledge requires infallible evidence. But we can reject the thesis that $E=K$, as we have done above, even while holding that knowledge requires infallible evidence.²⁰ Perceptual knowledge that x is F may require that one *perceive* x to be F – and this is something one can do even if one does not know that Fx. (There is a difference here between perceiving that x is F and perceiving x to be F. Both are sufficient for x to be F, but only the former is sufficient for *knowing* x to be F.) Testimonial knowledge that x is F may require that one be *informed* of x's being F – and this is something that can happen to one even if one refuses to believe that Fx, and so even if one fails to know that Fx. (Again, there is a difference here between being informed of x's being F and learning that x is F. I take it that both suffice for x to be F²¹, but only the latter suffices for knowing x to be F.) Inductive knowledge that A's are B's may require that one be *informed* of A's being B's – and this is something that can happen to one even if one refuses to believe that A's are B's, and so even if fails to know that A's are B's. It's possible for one perceive x to be F, or be informed of x's being F, or of A's being B's, without knowing that Fx, or that A's are B's. But it's not possible for one to perceive x to be F, or to be informed of x's being F, unless x is in fact F. And it's impossible to be informed of A's being B's without A's in fact being B's.

Some philosophers may worry that such truths as that one is informed of x's being F, or of A's being B's, cannot be in one's evidence set, for one's evidence set includes only those truths that one can know to be true solely by reflection (either reflection on the

contents of one's own mind, or *a priori* reflection). Now it is true that sometimes such truths are not in one's evidence. For instance, when Henry, in barn façade country, sees a barn to be in front of him, the fact that he sees a barn to be in front of him is not in his evidence set. When you read a newspaper story that informs you of the assassination of a famous official, the fact that you have been so informed is not in your evidence set if, say, later newspapers retract that story. To be informed of x's being F is one thing – to have the fact that one is so informed in one's evidence set is a further thing. Only if one is entitled to take it for granted in inquiry that one has been so informed does the fact of one's having been so informed make its way into one's evidence set.

But, according to the objection voiced above, such facts as one's having been informed of x's being F never make their way into one's evidence set, and this is because one cannot know such facts by reflection alone. Now, even supposing it is true that one's evidence set can include only those truths that one can know to be true by reflection alone (and I would need to hear why we should accept it) why cannot the present view accommodate this constraint on what can be in one's evidence set? Why cannot one know, by reflection alone, that one sees the table to be red? I have elsewhere argued that, pending further illumination concerning what it is to know something by reflection alone, there is no good answer to this question.²²

Knowledge, on the present view, is belief that is properly based on infallible evidence (indeed, on evidence that can be known – perhaps upon reflection alone – to be infallible). So it follows that knowledge is *also*, as the IJTB theory says, justified true belief that is indefeasible by any further evidence: if S knows that p on the basis of evidence e, then e is infallible evidence for p, and so the conjunction of e with any other

true proposition is also infallible evidence for p, and therefore constitutes a justification for S to believe that p. Of course there may be examples in which S is *justified* in believing that p on the basis of evidence which can be expanded into something that is not a justification for S to believe that p – justification itself may be defeasible. But knowledge is not defeasible, according to the IJTB theory. The theory that I have been calling the IJTB theory – the theory that says that knowledge is indefeasible justified true belief – is just a first approximation to the theory that I would like, in the end, to defend: the *infallibly* justified true belief theory of knowledge. For convenience, I will henceforth use the expression “IJTB theory” to refer to the latter theory of knowledge, the theory that expresses precisely what I hold, rather than the earlier approximation to it.

If I want to defend a theory according to which knowledge is infallibly justified true belief, then why did I make a detour through indefeasibility? Indefeasibility is an important feature of knowledge, because it is the indefeasibility of knowledge that explains why knowledge, and knowledge ascription, is important to us. Recall that we wanted a theory of knowledge to tell us what knowledge is, and we wanted to know what knowledge is in order to be able to explain how it is that knowledge, and knowledge ascriptions, can play the role that they play in our lives. We’re now in a position to provide at least some of the latter explanatory story.

IV. Why it matters what you know

Knowledge and knowledge ascription play many different roles in our lives. A complete account of knowledge would help us to understand how they manage to play these various roles. But a complete account of knowledge would have to say something

about the nature of the justification required for knowledge, the nature of the basing relation required for knowledge, and the nature of belief and truth. In this paper, I have said almost nothing about these complicated issues, and I will not say anything more about them in this section. I have focused exclusively on the anti-Gettier condition on knowledge, the indefeasibility condition. Can our discussion of this indefeasibility condition help us to understand what roles knowledge and knowledge ascription play in our lives, and how it is that they play those roles?

I believe it can, and in order to substantiate this, I will focus on just one example in this paper. According to the IJTB theory, knowledge is indefeasible, in the following sense: so long as the agent who knows that *p* retains all of her current evidence, her acquisition of new evidence will never render her justified in disbelieving that *p*. But then *so long as she retains her current evidence*, she will never be justified in disbelieving that *p*. Now suppose that such an agent adopts a dogmatic attitude towards *p*: she treats *p* as a fixed or settled point in inquiry, and so when she is confronted with recalcitrant evidence that might be thought to tell against *p*, she revises some belief other than *p*. What epistemic cost would such an agent incur? Of course, if the agent doesn't know that *p*, but only thinks that she knows that *p*, then she may end up incurring the epistemic cost of never finding out that she was wrong about *p*, or at least of holding beliefs that she is not justified in holding. But, by hypothesis, our agent knows that *p*, and retains all of her evidence for *p*, so, at least according to the IJTB theory, these are not costs that she could incur. Indeed, I cannot think of any cost that such an agent might incur by being dogmatic about *p*. If there is indeed no such cost, then knowing that *p* entitles one to be dogmatic about *p*, at least so long as one retains one's evidence for *p*.

To be dogmatic about p is not the same as being highly confident of p, or even rationally highly confident of p. Dogmatism need not involve a high degree of confidence or of rational confidence. One may be dogmatic about p but not about q, even if one is actually and/or rationally more confident of q than of p. Rather, dogmatism about p is a disposition to treat further evidence as not calling p into question. If I believe that p, and I am disposed to treat all further evidence as not calling p into question, then I am dogmatic about p. So long as I know that p, and I retain the evidence on the basis of which I possess this knowledge, rationality entitles me to be dogmatic about p, because there is no epistemic cost that I might incur by continuing to believe that p no matter what new evidence may come my way.

Many philosophers think that it is generally unreasonable to be dogmatic, even about those propositions that we know to be true. Witness the popular reaction to the dogmatism paradox that was formulated by Kripke, and first presented in print by Harman. Harman states the argument for the allegedly paradoxical conclusion as follows:

“If I know that h is true, I know that any evidence against h is evidence against something that is true; so I know that such evidence is misleading. But I should disregard evidence that I know is misleading. So, once I know that h is true, I am in a position to disregard any future evidence that seems to tell against h.” (Harman 1973, 148.)²³

After stating this argument for dogmatism, Harman (and several other philosophers following Harman) criticizes the argument by claiming that gaining new evidence can make one lose one’s knowledge. Now, the proponent of IJTB can agree that gaining new evidence can make one lose one’s knowledge – if this connection is understood as a

causal connection. Gaining new evidence can cause one to lose one's old evidence for p, or can cause one to lose one's belief that p, and either way it can cause one to lose one's knowledge that p. But, according to IJTB, gaining new evidence cannot, all by itself and apart from any effects it may have, cause one to lose one's knowledge or one's justification for believing that p. So while Harman and other anti-dogmatist philosophers are right to say that gaining new evidence can make one lose one's knowledge, this is compatible with the thesis of the indefeasibility of knowledge that's being advanced here. If it is generally unreasonable to retain one's beliefs in the face of apparently contrary evidence, that is because gaining such evidence can shake one's confidence, or cause one to lose some of one's old evidence, and either way it can rob one of knowledge. But it cannot rob one of knowledge simply by tipping the evidentiary scales, without depriving one of any old evidence.

So what role does knowledge, and knowledge ascription, play in our lives? The proponent of IJTB is in a position to answer this question as follows. Knowledge is important to us because it entitles us to hold certain points fixed in inquiry. Knowledge ascription is important to us because it records those points that we are entitled to *hold fixed* in inquiry. (What we are entitled to hold fixed in inquiry is not the same as what we are entitled to take for granted in inquiry, even though there is overlap in the extension of the two categories. For instance, some of what I am entitled to hold fixed I am not entitled to take for granted, because I must be able to defend it against challenges.) If we didn't have knowledge, then we really would be in a predicament no better than the one that Neurath described: having to rebuild the raft of our opinions while we are afloat on it. Those of us who prefer to stay safely on dry land while rebuilding our raft should be

grateful therefore that we do actually know a lot, and we do not have to settle for merely justified true beliefs.²⁴

Works Cited

- Ayer, A.J. 1956. *The Problem of Knowledge*. St. Martin's Press: New York.
- Chisholm, Roderick. 1957. *Perceiving*. Cornell University Press: Ithaca, NY.
- . 1977. *Theory of Knowledge, Second Edition*. Prentice-Hall: Englewood Cliffs, NJ.
- Dretske, Fred. 1971. "Conclusive Reasons." *Australasian Journal of Philosophy* 49: 1 – 22.
- . 1981. *Knowledge and the Flow of Information*. MIT Press: Cambridge, MA.
- French, P., Uehling, T., and Wettstein, H., eds. 1980. *Midwest Studies in Philosophy, Volume V: Studies in Epistemology*. University of Minnesota Press: Minneapolis, MN.
- Ginet, Carl. 1980. "Knowing Less by Knowing More" in French, Uehling, and Wettstein 1980.
- Goldman, Alvin. 1967. "A Causal Theory of Knowing." *Journal of Philosophy* 64: 355 – 72.
- . 1976. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73: 771 – 91.
- Harman, Gilbert. 1968. "Knowledge, Inference, and Explanation." *American Philosophical Quarterly* 5: 164 – 73.
- . 1970. "Knowledge, Reasons, and Causes." *Journal of Philosophy* 67: 841 – 55.
- . 1973. *Thought*. Princeton University Press: Princeton, NJ.
- Hawthorne, John. 2003. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Klein, Peter. 1971. "A Proposed Definition of Propositional Knowledge." *Journal of Philosophy* 68: 471 – 82.
- Lehrer, K. and Paxson, T. 1969. "Knowledge: Undefeated Justified True Belief." *Journal of Philosophy* 66: 225 – 37.
- Lewis, David. 1973. *Counterfactuals*. Harvard University Press: Cambridge, MA.
- Neta, Ram. 2004. "Skepticism, Abductivism, and the Explanatory Gap." *Philosophical Issues* 10: 296 – 325.
- Nozick, Robert. 1981. *Philosophical Explanations*. Belknap Press: Cambridge, MA.

- Pritchard, Duncan. 2004. *Epistemic Luck*. Oxford University Press: Oxford.
- Rescher, Nicholas, ed. 1968. *Studies in Logical Theory*. Blackwell: Oxford.
- Rieber, Steven. 1999. "Knowledge and Contrastive Explanation." *Nous* 32: 189 – 204.
- Sosa, Ernest. 1999. "How to Defeat Opposition to Moore." *Philosophical Perspectives* 13: 141 – 54.
- . 2004. "Relevant Alternatives, Contextualism Included." *Philosophical Studies* 119: 35 – 65.
- Stalnaker, Robert. 1968. "A Theory of Conditionals" in Rescher 1968.
- Unger, Peter. 1975. *Ignorance*. Clarendon Press: Oxford.
- Weinberg, Stich, and Nichols. 2001. "Normativity and Epistemic Intuitions." *Philosophical Topics* 29: 429 – 60.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press: Oxford.

¹ Throughout this paper, I use the term "knowledge" to refer solely to factual knowledge, i.e., the kind of knowledge that is ascribed using assertions of the form "S knows that p", where "S" is a schematic variable that ranges over persons and "p" is a schematic variable that ranges over propositions. Factual knowledge is also, I believe, ascribed using assertions of the form "S knows wh— ...".

² Of course, we could reject Williamson's equation $E=K$, but still accept the view that p is included in S's evidence at time t IFF S knows that p at t. We could do this by saying (contra Williamson) that what it is to be in S's evidence is something other than simply to be known by S, but nonetheless, although it is not part of the very essence of being S's evidence that something is known by S, the biconditional is still metaphysically necessary. If Williamson were simply to hold that the biconditional is metaphysically necessary, without holding $E=K$, then he could still accept the account of knowledge that I offer below as informative. For reasons that I offer below and elsewhere, I do not even accept that the biconditional is true, let alone metaphysically necessary.

³ See, e.g., Ayer 1956 and Chisholm 1957.

⁴ Some people deny that this was discovered. For instance, Weinberg, Stich and Nichols 2001 claim that Gettier intuitions are culturally variable, and that this shows that they do not carry probative weight in epistemology. I am happy to grant that such intuitions are culturally variable, but I do not see why this suggests that they do not carry probative weight. The Mediterranean feels cold to Nigerians and hot to Eskimos – does this suggest that thermal sensation is not an epistemologically respectable guide to real thermal phenomena? It shows only that, when relying on thermal sensations, we need to discount the possible distorting influence of environment. If Gettier intuitions are the result of some such distorting influence, then of course we should discount them. But why should we think that they are the result of some such distorting influence? Simply because some people don't share them? Some people (e.g., non-English speakers) don't share our intuitions about the grammaticality of English sentences, and some people (e.g., xenophobes) don't share our intuitions about the irrelevance of race or ethnicity to moral status. But that has no tendency to show that our intuitions about the grammaticality of English sentences, or our intuitions about the irrelevance of race or ethnicity to moral status, are the result of distorting influence.

⁵ Harman 1968, 1970, 1973.

⁶ Goldman 1967 defends a causal theory of at least some kinds of empirical knowledge. The counterexample to it that is presented in the text was first published in Goldman 1976.

⁷ See, e.g., Rieber 1999.

⁸ Dretske 1971 propounds a version of the Sensitivity Theory, as does Goldman 1976. Nozick 1981 propounds a theory slightly stronger than the Sensitivity Theory: his theory adds a clause to the effect that, in all nearby p-worlds in which the believer forms a belief as to whether p, using the same method M, the believer continues to believe that p. But though Nozick's version is stronger, it still doesn't handle the counterexample offered in the text.

⁹ Lewis 1973, Stalnaker 1968.

¹⁰ Though Williamson 2000 and Sosa 1999 both defend a safety condition on knowledge, neither would accept what I am calling the Safety Theory. Pritchard 2004 argues that Safety is the anti-Gettier condition on knowledge. But if knowledge is justified, true unGettiered belief, and if Safety is the anti-Gettier condition, then knowledge is justified true safe belief. Since (I am assuming here) knowledge is justified true unGettiered belief, and since (I argue in the text) knowledge is not justified true safe belief, I conclude that Safety is not the anti-Gettier condition on knowledge.

¹¹ Lehrer and Paxson 1969 mention this theory but do not endorse it. I don't know of any philosophers who endorse it, but I mention it only in order to distinguish it from my IJT theory, which is superficially similar.

¹² See Shope 1983 for an encyclopedic review of counterexamples of this form.

¹³ Chisholm 1977.

¹⁴ Williamson 2000, 219.

¹⁵ I defend this view in my manuscript "What Evidence Do You Have?".

¹⁶ Let E be my current evidence set. Can E provide me with conclusive evidence for any proposition? It can do so only if it provides me with maximal justification for believing that proposition. But can't my justification for any proposition be improved? In fact, do I not improve my justification – however strong it may already be – by not merely relying upon my evidence, but also defending my evidence against challenges? I think the answer to this question is "no – not always". For instance, I do not improve my justification for believing, say, that I have hands, each time I get my eyes checked. Many epistemologists have views that commit them to disagreeing with me about this last claim. But on this issue, it seems to me, their views have a wildly implausible consequence.

¹⁷ For more, see my "What is Your Evidence?"

¹⁸ For that matter, it is also not a good reason to reject Jeffrey's rule of conditionalization. It is a consequence of Jeffrey's rule that, if $\text{Prob}(p/e) = 1$, then new evidence cannot, *by itself*, reduce the degree of credence one should lend p. But it is not a consequence of Jeffrey's rule that new evidence cannot cause one to lose one's old evidence, and thereby reduce the degree of credence one should lend p.

¹⁹ If fallible evidence cannot give us knowledge, then can the inductive argument have given us knowledge that its conclusion is true? Yes. Inductive arguments can give us infallible evidence for their conclusion, as I argue below in the text.

²⁰ In "McDowell and the New Evil Genius", Duncan Pritchard and I interpret John McDowell as holding this view.

²¹ Bill Lycan reports hearing a use of "inform" on which being informed that x is F is compatible with x's not being F. I confess that I do not hear such a use of "inform", but perhaps I am being deaf. If so, then I should find a better word to denote the factive relation at issue here. Perhaps English does not supply such a word. If such a word does not exist, epistemology will have to invent it.

²² I discuss this complicated issue more fully in my "Refutation of Internalist Fallibilism".

²³ Ginet 1980 notes that Harman's argument, as stated, is enthymematic, he reformulates it thus:

For any propositions h and f:

(Premise 1) If (a) I know that h is true, then (b) I know that, if in the future I come to know f to be true, my then counting f as evidence against h would be to count f as evidence against a truth;

(Premise 2) If (b) I know that if in the future I come to know f to be true then my counting f as evidence against h would be to count f as evidence against a truth, and I do not forget this knowledge, then if in the future I come to know that f is true then at that time I know that my counting f as evidence against h would be to count f as evidence against a truth;

(Premise 3) If (d) I know that f is true and that my counting f as evidence against h would be to count f as evidence against a truth, then I am entitled not to count f as evidence against h;

(Conclusion) Therefore, if (a) I know that h is true, and I do not forget this knowledge, then (e) if in the future I come to know that f is true, then at that time I am entitled not to count f as evidence against h.

Ginet rejects Premise 2 of the argument, so formulated, and he plausibly takes Harman to be best understood as implicitly rejecting Premise 2 as well.

²⁴ For comments and discussion, I am grateful to Bill Lycan, John Roberts, Dylan Sabo, and other members of an audience at the University of North Carolina, Chapel Hill. I am also grateful, as always, to Duncan Pritchard and to Jonathan Schaffer.