

# Evidence Neutrality

Brian Weatherson \*

April 13, 2009

## 1 Psychologising Evidence

Chapter 7 of Timothy Williamson's *The Philosophy of Philosophy* is an extended argument against psychologising evidence in philosophy. Before we can evaluate those arguments, it would be useful to get clear on just what it is to psychologise evidence. In this post I'll say a bit about what that amounts to, and in the next post look a bit more carefully at Williamson's text to see just what position he is attributing to his opponent.

In some ways the debate Williamson is contributing to is among the oldest in modern philosophy. Consider the following two positions about perceptual evidence, each of which has found many partisans over the last few centuries.

- Perceptual Evidence is Psychological. My perceptual evidence consists in facts about the psychological states I am in when undergoing a perceptual experience. So, for instance, my perceptual evidence might include that I'm visually representing that there is a table in front of me.
- Perceptual Evidence is External. My perceptual evidence consists in facts that I perceive. So, for instance, my perceptual evidence might include that there indeed is a table in front of me.

The psychological theory has a number of advantages. It can explain how people having illusory perceptions can get the same kind of evidence (albeit of lower quality) as people having veridical experiences. It arguably staves off certain kinds of doubts about our evidence, at least to the extent that we have privileged access to our psychological states. It explains the fact (if it is a fact) that when we get evidence in favour of some proposition  $p$  about the external world, we generally know what kind of evidence we have. It is unusual, that is, to get evidence that  $p$ , but not know whether that is visual evidence, or tactile evidence, or testimonial evidence, or whatever. If the evidence for  $p$  just is the visual or tactile or testimonial experience, that is easily explained. And it offers the prospect of an easy theory of evidence possession; a point I'll return to below.

But there's one big cost of the psychological theory: it seems to promote scepticism. There is a long tradition, starting in the modern period with Descartes, of proponents of the psychological view wondering how to get from psychological evidence to knowledge of the external world. And there is another long tradition, culminating at the present with Williamson, of opponents of the psychological view using this worry as a reason to start with evidence in the external world, and avoid this sceptical doubt.

The debate here is not confined to perception. We can have a similar debate in testimony. Imagine I am told that  $p$  by a trusted friend. I now have some evidence for  $p$ . What is it? One answer, similar in spirit to the psychological

---

\*Thanks to Jonathan Ichikawa, Ishani Maitra, Crispin Wright and many others. References etc incomplete.

answer above, is that I've been told that  $p$ . Another answer, similar in spirit to the external answer, is  $p$  itself. The latter answer might be favoured by a theorist of testimony who thinks that when I get testimony from a trustworthy source, I simply receive the warrant they have for believing  $p$ . (The two answers here aren't quite equivalent to the positions known as reductionism and anti-reductionism in the epistemology of testimony. Someone might be an anti-reductionist and hold that the telling, rather than what's told, is the evidence, by holding that we don't need any extra grounds to infer, on the basis of that evidence, that  $p$ . I'll say more about such inferential rules in later posts.)

Both answers here are possible, but it is much more plausible to take the evidence to be the telling rather than what's told. So we can use that as a relatively clear example of what happens when we take evidence to be something that supports an external world proposition  $p$ , rather than  $p$  itself. One consequence is that in reporting inferences, we can replace testimonial knowledge with knowledge that the testimony was made without making the inference worse. So imagine we know that if Celtic win today, they're champions, and we're told by a trusted friend that Celtic did indeed win. Then we might make either of the following inferences.

An inference from facts about football

1. If Celtic won, they are champions. 2. Celtic won. 3. So, Celtic are champions.

An inference from facts about testimony

1. If Celtic won, they are champions. 2. My friend said that Celtic won. 3. So, Celtic are champions.

The first is valid, while the second is not. But we are interested here in inferences, not implications, so that's no disqualifying mark against the second inference. For the second has a virtue not shared by the first, namely that its premises are more secure. So it looks like the two inferences are equally good. And that suggests that the second inference really is just making explicit the inference that's underlying the first.

We've now said enough to set up the interesting debate about philosophical evidence. Often we say things like Intuitively, that's a cause of that, or Intuitively, that's wrong and so on in philosophy. What kind of evidence are we appealing to here?

- Psychological States. Our evidence is that we have the intuition that such and such is a cause, or is wrong, or so on. In other words, our evidence is a certain psychological state.
- External Facts. Our evidence is not the intuition itself, but its content. So our evidence is that such and such is a cause, or is wrong, or whatever. When we say that this is intuitive, we are perhaps reporting how we have that evidence, we are not describing the evidence we have.

The philosopher who psychologises philosophical evidence is the one who says that our evidence is the first kind, the psychological states. So far it seems clear that this is the kind of view that Williamson has in mind. But we might wonder how far this goes. An extreme version of the view that philosophical evidence is psychological is that we could always replace  $p$  in a philosophical argument with intuitively  $p$ , without loss of argumentative force. Remember in the testimony case we replaced a fact we were told about with the fact of the telling and didn't seem to make our argument any worse; the same goes on here.

That position is implausible in the extreme. It doesn't need any fancy argument about the nature of evidence to see that it's wrong, it simply mistakes the nature and scope of philosophy. To see this, consider an argument loosely based on Peter Singer's *Animal Liberation*.

1. Eating meat is morally permissible only if it is necessary for a healthy diet 2. Eating meat is not necessary for a healthy diet 3. So, eating meat is not morally permissible

I don't know whether that's a sound argument, but it is a recognisably philosophical argument, and an interesting one. The following is not, however, a particularly good or interesting argument.

1. Eating meat is morally permissible only if it is necessary for a healthy diet 2. Intuitively, eating meat is not necessary for a healthy diet 3. So, eating meat is not morally permissible

If Singer had presented that argument, he wouldn't have been listened to. The crucial evidence for premise 2 is from nutrition science, not intuition. So it's not always true that we can replace  $p$  with Intuitively,  $p$  in philosophical arguments without loss of argumentative force. So it's not always true that making this replacement merely makes more explicit our underlying reasoning. So it's not always true that our philosophical evidence consists in intuitions.

All that shows is that not all philosophical evidence consists of intuitions. It doesn't show that, for example, no philosophical evidence consists of intuitions. One way to psychologise evidence' is to say that some of the time, or perhaps even frequently, intuitions are a part of our philosophical evidence. If this is plausible, then the interesting question seems to be scalar rather than on/off. It isn't whether evidence is always intuitions or always something else; it's how often does evidence consist of intuitions?

There's a second scalar question we might wonder about, one triggered by the observation that the two positions on evidence we've been taking are not strictly speaking incompatible. We could hold, indeed a sensible proponent of the external view probably will hold, that psychological states are among our evidence, even when our evidence also consists in external facts. So in the perceptual case, my evidence might include both that there is a table in front of me, and that I'm forming a visual representation of the table.

When this is true, there are a number of questions we might start to think about concerning the balance between the psychological evidence and the external evidence. For purposes of these posts, what I'll be interested in is the following question. In such cases, how well does the purely psychological evidence support the conclusion? And in particular, how much better supported is the conclusion by external evidence than it would be by psychological evidence alone? If the answers to those questions are quite a bit and not a lot, then the psychological view seems to have turned out to be correct in crucial respects.

Note this doesn't mean that the psychological states have to do a lot of work. Consider again the position that I get both external and psychological evidence when I look at a table. One might hold (a) the psychological evidence I get provides very strong grounds for believing that there's a table, while (b) I don't actually need that evidence, since I have even stronger evidence available, namely that there is a table. That still feels like a case where the person who takes philosophical evidence to be psychological isn't badly mistaken, even if in fact the psychological evidence in the case doesn't need to do any work. If (a) were false, and the psychological evidence provided very little by way of justification, then the psychological view does feel like it is mistaken, even if strictly speaking, my psychological states are (weak) evidence that I actually have.

So here's a somewhat more precise way to say what it is to psychologise evidence in philosophy: it's to hold that in a large percentage of philosophical argument, a large amount of the evidence consists of intuitions, and by that we mean that intuitions alone can bear a lot of the weight of the philosophical argument even in the absence of non-psychological evidence. In the next post I'll work through some of the things Williamson says to see how close that position is to his target, and then in the posts to follow, I'll sketch a defence of such a position.

## 2 Williamson on Evidence

In the previous post I mentioned that Williamson clearly opposes in chapter 7 a broadly psychological conception of philosophical evidence. But it isn't exactly clear just what his target is. At times he seems to be arguing against psychological evidence ever being philosophically worthwhile. For example, consider the following batch of quotes.

For now I face the challenge of arguing from a psychological premise, that I believe or we are inclined to believe the Gettier proposition, to an epistemological conclusion, the Gettier proposition itself. That gap is not easily bridged.

Since psychological evidence has no obvious bearing on many philosophical issues, judgment scepticism is also encouraged in ways that do not depend on the consequence fallacy.

In explaining why we have intuitions, analytic philosophy has a preference for explanations that make those intuitions true over explanations that make them untrue, but the justification for that preference remains unclear

In those quotes his opponent seems to hold the relatively weak view that psychological evidence can (sometimes) be useful evidence for philosophical conclusion. But other times he seems to take his opponent to be the person who holds the much stronger view that only intuitions are evidence. For instance, he says

One result [of EN] is the uneasy conception many contemporary analytic philosophers have of their own methodology. They think that, in philosophy, ultimately our evidence consists only of intuitions.

I think it's hard to believe that's really a widespread view in philosophy. Does Singer's argument for vegetarianism rest (even ultimately) on intuitions about the nutritional value of a vegetarian diet? Does the well-known argument from special relativity against presentism rest on intuitions about whether special relativity is true? Nevertheless, Williamson does attribute it to many (unnamed) philosophers. And yet some of Williamson's arguments seem directed particularly against this position. For instance, he says

Taken far enough, the psychologisation of philosophical method becomes self-defeating

And he cites approvingly Joel Pust's conclusion that it is self-defeating to hold that

Aside from propositions describing the occurrence of her judgements, S is justified in believing only those propositions which are part of the best explanation of S's making the judgements that she makes

Probably Pust and Williamson are right here, but it hardly tells against anything but a strawman version of the psychological view of evidence. Finally, Williamson objects to a version of Reflective Equilibrium that just attempts to get our intuitions into equilibrium with the following argument.

The reflective equilibrium account, as usually understood, already assigns a proto-evidential role to at least one kind of non-psychological fact. For it treats philosophers as relying on logical relations between theories and intuitions, in particular their consistency and inconsistency.

The theme again is that we need some evidence other than intuitions, something that should be common ground. (For reasons I'll suggest in the next post, I'm not sure this is a good argument for that conclusion though.)

We haven't got very far by trying to characterise what Williamson's opponent says. Perhaps it is better to look at his positive proposal for what is evidence in philosophy. We get one statement of what that positive conclusion.

Our evidence in philosophy consists of facts, most of them non-psychological, to which we have appropriate epistemic access.

That, in conjunction with the quotes above, suggests he is defending the following three theses.

1. Not all philosophical evidence is psychological. 2. Having  $p$  be part of your evidence requires appropriate epistemic access to  $p$ . 3. The intuition that  $p$ , or the fact that one has that intuition, is weak evidence, perhaps no evidence at all, that  $p$ .

A position that denied all three of those would clearly be among the targets of Williamson's chapter. But that would be crazy, since (1) is obviously true. But a position like the one I sketched at the end of the last post, which was neutral on (2) and denied (3), would still seem to be at odds with the bulk of what Williamson says, and I think is meant to be among the positions ruled out by the considerations he raises.

Now such a position does not seem especially related to the sceptical positions that are the targets of sections 3 and 4 of Williamson's chapter 7. But that's as things should be. The question of what our evidence is doesn't immediately settle the question of what knowledge we have. Taking perceptual evidence to be psychological might be a precursor to defending external world scepticism. But it might also be a precursor to adopting indirect realism. Of course Williamson, by accepting knowledge as evidence, has effectively ruled out classical forms of indirect realism, where we know about the world on the basis of purely phenomenal evidence. But that shouldn't be presupposed here I think. A position that holds that psychological states, or facts about them, are often crucial evidence for us is opposed to the core doctrines of Williamson's chapter, even if it is also opposed to some of his other opponents.

### 3 Evidence Neutrality and Rules of Evidence

In chapter 7 of *The Philosophy of Philosophy*, Timothy Williamson argues against the idea that intuitions are a key part of philosophical evidence. Part of his argument is indirect. He thinks the motivation for taking intuitions to be central comes from accepting a principle he calls Evidence Neutrality, and that that principle is false. I rather suspect that isn't the best reason to take intuitions to be philosophical evidence, but we'll set that aside here. What we'll focus on here is whether Evidence Neutrality is true. Here is the initial statement of Evidence Neutrality.

Although the complete elimination of accidental mistakes and confusions is virtually impossible, we might hope that whether a proposition constitutes evidence is in principle uncontentiously decidable, in the sense that a community of inquirers can always in principle achieve common knowledge as to whether any given proposition constitutes evidence for the inquiry.

It seems to me that this is ambiguous between two readings.

- The weaker reading is that it is decidable, by consensus, which propositions are, in principle, evidentially relevant to an inquiry as to whether  $p$ .
- The stronger reading is that that is true, and it is also decidable, by consensus, in which epistemic direction each piece of evidence points .

I'm going to argue that one of the arguments against Evidence Neutrality, what we might call the argument from extremists, does not tell against the weaker version. I think (though this claim will eventually need defending) that if either version of Evidence Neutrality is metaphilosophically interesting, then the weaker version is interesting. So perhaps that's all that we need to defend.

Evidence Neutrality (hereafter, EN) is a kind of dialectical conception of evidence (hereafter, DCE). What our evidence is just is what our interlocutors will allow as evidence. On the stronger reading, it is what our interlocutors will take to be evidence for our conclusions. On the weaker reading, it is what they'll allow as evidence, though they may say one particular piece of evidence, a piece we take to be crucial, is not very strong. And the big question here is whether we should think of evidence dialectically.

It's certainly true that evidence that is accepted by our interlocutors will be more persuasive in convincing interlocutors. But that's no argument, at least no immediate argument, for a DCE. It might be that we have quite a lot of evidence that tells whether *p*, and our interlocutors are just mistaken about this. (Everyone makes mistakes.) Relatedly, some people may simply fail to be persuaded by arguments that are rationally persuasive. So we shouldn't simply confuse which evidence is dialectically effective with which evidence is genuinely good. If we want to defend a DCE, we'll have to argue for it more carefully than that.

The key point of the last paragraph is that some people will fail to be persuaded by genuinely good arguments. That suggests a problem; couldn't we have evidence against a position, but just not evidence accepted by the partisans of that position? The simplest examples of this will be positions whose partisans are hostile to the very idea that evidence can tell in favour of anything at all. Here is how Williamson converts such examples to arguments against EN.

Some scepticism, like scepticism about reason, is so radical that it leaves too little unchallenged for what remains as shared evidence to be an appropriate basis for evaluating the claims under challenge.

The point here is not a new one. David Lewis makes a similar observation in *Logic for Equivocators*.

The radical case for relevance [i.e. dialethism] should be dismissed just because the hypothesis it requires us to entertain is inconsistent. That may seem dogmatic. And it is: I am affirming the very thesis that Routley and Priest have called into question and contrary to the rules of debate I decline to defend it. Further, I concede that it is indefensible against their challenge. They have called so much into question that I have no foothold on undisputed ground. So much the worse for the demand that philosophers always must be ready to defend their theses under the rules of debate.

The point Williamson and Lewis make is clear enough. There are certain radical views that (a) we know to be mistaken, but (b) the nature of the position is such that it has, by its own lights, defences against the actual grounds for our knowledge that it is mistaken. Of course its lights are bad lights; our reasons are good reasons. But such positions have partisans. (This is clearer in Lewis's case than in Williamson's.) If our only evidence is the evidence they'll let us share, we won't have evidence against these positions. And that might suggest we don't really know the positions are mistaken, contrary to assumption. (There is a fairly strong evidentialist assumption being made here, namely that if we don't have evidence against such positions, we don't know they are mistaken. It's worth thinking through whether that assumption is right, but I won't do it here.)

I think, however, that this point goes by too fast. Remember that EN and DCE are claims about evidence. They aren't claims about what we can do with evidence. To see the importance of this distinction, it's worth recalling

Lewis Carroll's fable of Achilles and the Tortoise. (The points to follow are perhaps familiar from recent work of Paul Boghossian and Crispin Wright. And I'm indebted here to discussions with Crispin. But note that I'm expressly not committing myself to Boghossian's views about the meanings of the logical connectives.)

Achilles knows  $p$ , and  $p \rightarrow q$ . He wants to infer  $q$ . The tortoise says, wait a second, are you sure that's a good inference? Achilles says he is sure. He's sure, he says, that  $(p \wedge (p \rightarrow q)) \rightarrow q$ . The tortoise thinks for a second, and then says that that does sound right. Let's have that as another premise he says. Achilles happily agrees, and then proceeds to infer  $q$ . The tortoise is still not sure. He wants to know how Achilles is drawing that conclusion. Achilles says he's sure that if  $(p \wedge (p \rightarrow q)) \wedge (p \wedge (p \rightarrow q)) \rightarrow q$  then  $q$ . The tortoise agrees that looks true, and says it seems like a pretty good premise to have. Achilles tries again to infer  $q$ , and the tortoise is again worried about why he's drawing that conclusion. The story continues for a surprising while, with Achilles adding more and more premises, and seemingly getting no closer to overcoming the Tortoise's worries.

There's a mundane lesson to be drawn from that, and an exciting lesson. The mundane lesson is that there is a distinction between premises and rules. Indeed, in every axiomatic formal system, we are given both axioms and rules to generate theorems from old axioms/theorems. In some simple systems the only rule might be modus ponens, the rule that Achilles was looking for. In other systems we might need a rule like necessitation, or universal-introduction. But we always need something more than just axioms.

The exciting lesson is that rules aren't the kind of things that stand in need of rational justification. They are, to put it perhaps in Wittgensteinian terms, things that justify, rather than things that are justified. Here is how we might draw that conclusion. We can imagine the tortoise not as an unhelpful interlocutor, but as our own nagging doubts. Our own inner Descartes, if you like. If the rules have justifications, then we should be able to give them. And if we give them, we can add them as extra premises from which we reason. But this is the key mistake Achilles makes. At some point we need to stop adding premises, and start doing something with the premises. And that can't always be supported by reasons. For imagine it could. That is, imagine the rule that let us go from  $A$  to  $B$  could be supported by evidence  $E$ . Then we can still ask, what's the rule that lets us go from  $A$  and  $E$  to  $B$ ? Still we'll need a rule, and perhaps now we'll be out of evidence. At some point a jump needs to be made without evidence.

So I conclude rules don't need evidential justification. That's not to say that all rules are created equally. There are normative standards governing rules, even though they are not supported by evidence. This makes their status quite delicate. As I read him, Gilbert Ryle introduced the idea of knowledge how directly to address this problem. Following rules can't be simply propositional knowledge, because that leads to a regress. On the other hand, following rules is normatively, even rationally, evaluable. Ryle thought that if we recognise a category of know how, we can steer between these rocks; we can have something that's a kind of knowledge, the exercise of which can be rational or irrational, but which doesn't require evidence.

If it isn't required that we be able to justify our use of rules to ourselves, it doesn't seem like it should be required that we be able to justify them to our friends. And that in turn suggests that a dialectical conception of rules would be inappropriate. Who cares if our (rational) friends don't like the rules we're using? The only way we could make them like them is by offering reasons that our rules are good rules, and by hypothesis we don't even need to be able to articulate such reasons to ourselves. Perhaps we don't even need to have such reasons. So a dialectical conception of rules is bad, and more specifically, Rule Neutrality (understood along the same lines as Evidence Neutrality) is bad.

But note that once we ditch Rule Neutrality, we can respond to the extremists that Lewis and Williamson are worried about without sacrificing Evidence Neutrality. Here's my evidence that dialethism is false. If dialethism is

true, some contradiction is true. Taking that to be evidence doesn't violate Evidence Neutrality, because it's agreed on all sides. From that it follows, by a rule that I properly accept (i.e. *reductio*) that dialethism is false. Of course, the dialethists don't buy that rule. But that's not my problem, since I'm only committed to sharing evidence with them, not sharing rules. If I accepted the strong form of Evidence Neutrality, that might be a problem, because of course the dialethists don't think this is evidence against dialethism. On the weak form of Evidence Neutrality, that isn't a problem either.

It's a little trickier to respond to the reasons sceptic, but I think it can be done, especially if we think about induction. So imagine that I see a lot of Fs that are all Gs, and I see them in a lot of different places etc. I conclude that I have good reason to believe the next F I see will be G. This is a direct inference; there is no mediating premise. If you don't think so, try to imagine (a) what such a premise could be, and (b) how it could be justified? I think there aren't good answers to this question, or at least that any answer is less certain than I am in the conclusion. So my frequent observation of green emeralds is sufficient evidence to conclude that I have a reason to believe something, and hence that reasons scepticism is false.

Summing up, I think that Williamson here has run together two similar, but importantly distinct, principles: Evidence Neutrality and Rule Neutrality. I think he's right that if you accept both, you'll have thrown away all hope of a good response to certain positions to which there are good responses. So we shouldn't accept both of those principles. But if we accept that evidence is knowledge, as Williamson does, then we should think that all our evidence requires justification. And we shouldn't think that our rules do. Since the acceptability of our evidence/rules to our (rational) interlocutors is grounded in this need for justification, it seems that our reason to accept Evidence Neutrality is not a reason to accept Rule Neutrality. So Rule Neutrality must go. And when it does, the argument from extremism against Evidence Neutrality goes too.

There's a lot to say about rules, and I'll say a very little about it tomorrow.

NB: Ichikawa's objection that intuitively we do have lots of evidence that our opponents reject. The reply is that we shouldn't appeal to it. EN is a methodological principle, not an epistemological principle.

## 4 Rules Without Justification

In the previous post we argued that as well as evidence, we need a notion of an unjustified rule that takes us from evidence to conclusion. Some may think that this notion is too obscure, or at least philosophically disreputable to do the work it's put to. This section is then a discussion of rules that aims to increase their respectability. The idea is obviously not new; it traces back as least as far as Wittgenstein on rule-following, if not back to Carroll himself. And it has links to contemporary epistemology, wherever someone says that there are things we are entitled to assume without argument. But since it is doing so much work here, and since in conversation it has often been the most puzzling aspect of the argument to others, it may be worth saying just a little about what I'm taking rules to be here.

The argument in that section relied on Lewis Carroll's example involving *modus ponens*, and that might suggest that *modus ponens* is the kind of rule that we need. Indeed, it often seems that some philosophers think that it's the only rule we need. (Many philosophers have said that conditionalisation, which is really just a probabilistic form of *modus ponens*, is the only rule we need.) I think this example is misleading for four related reasons.

First, *modus ponens* is a rule of implication and what we're really looking for here are rules of inference. We're looking for rules that tell you what to do with evidence. And what you do with evidence is draw inferences from it.



These may not, in any interesting sense, be implications of the evidence.

Second, modus ponens is necessarily truth preserving, and this might be thought to be related to its acceptability. This seems like a mistake twice over to me. It's arguable (indeed I'll briefly argue for it presently) that we need some ampliative rules to explain the rationality of induction. And there are necessarily truth preserving rules that we cannot employ without justification. The rule that lets us infer  $p$  and oxygen has atomic number 8 from  $p$  is necessarily truth preserving, but not a rule we could freely employ without independent justification.

Third, modus ponens can be used in suppositional reasoning just as well as in regular reasoning. It would take us way too far afield to investigate this properly, but I suspect some of the rules we'll look at can't be properly be used in all suppositions. (Of course some rules of implication, at least in some formulations, also have restrictions on when they can be used; think of the restrictions on necessitation or universal-introduction.)

Fourth, modus ponens might (although this is controversial) be constitutively related to the meaning of the conditional. Perhaps, as some inferentialists believe, the meaning is determined by the acceptability of the rule. Perhaps the meaning directly implies that it is an acceptable rule. If either of those things are true, they aren't I think things we'd want to generalise to, say, rules for rational inductive inference. (There is obviously a lot to be said here, such as considering what might justify inferences related to logic if not meaning. Many recent papers by Boghossian, Williamson, Wright and others are relevant here. But I'm just going to avoid that issue for today, in large part because I do find it mysterious how to generalise much of that debate from issues about modus ponens to issues about, say, enumerative induction.)

Now all that is just to say what rules are not. Can we say what they are? As noted above, in some ways rules are very familiar. Any time any philosopher claims that we are warranted, or entitled, to, without justification, rely on the deliverances of some source, it seems they are proposing that the inference from The source says  $p$  to  $p$  is a good rule. So we can take dogmatists about perception (e.g. Pryor's *The Sceptic and the Dogmatist*) as endorsing a rule that lets us move from Appears that  $p$  to  $p$  without antecedent justification. And some anti-reductionists about testimony seem to hold that the rule which licences the inference from I'm told that  $p$  to  $p$  is a good one. So in some ways this isn't a new idea, it's just a way of framing an old idea. There are, however, two reasons we might think that this is a good framing. First, it lets us ask some relatively precise questions about the statement of the rules. Second, by thinking about rules as a class, we can formulate restrictions on what could be a rule.

It is harder than one might like to actually state rules that we can or do use. Clearly it isn't a rule that we can, in every case, infer from Appears that  $p$  to  $p$ ; there are illusions, some of which we know about. Nor is it a rule that we can, in every case, infer from I have observed many Fs in widespread conditions, and they have all been Gs to The next F I observe will be G; there are gruesome predicates. We could try to incorporate the exceptions into the rule, but a quick glance at how one might do this reveals that it isn't too attractive, at least as long as we want to have non-trivial rules. Probably we need something different.

In the case of appearance, there is a natural move to make. Instead of a single rule with all sorts of qualifications, we might try to defend the following two unqualified rules.

(A) If it appears to you that  $p$ , then you have a reason to believe that  $p$ . (ND) If you have a reason to believe that  $p$ , and this reason isn't defeated, believe that  $p$

The conjunction of (A) and (ND) entails Pryor's dogmatism, although that conjunction is considerably stronger than dogmatism. For example, it entails that appearances that you have reason to believe are deceptive provide some reasons; dogmatism as such is silent on such appearances. But both rules seem at least defensible, even given the

existence of illusions. (Of course, if they are good rules, they don't need a defence. But we can reasonably expect that a good rule won't be such that we can have good reason to believe it is bad, and (A) and (ND) pass that rule.)

The case of induction is a little trickier, because as soon as we think about the rule a hard question arises. We can only project non-gruesome predicates. Does that mean (a) that the non-gruesomeness of the projected predicates should be an input to inferences licenced by the rule, or (b) that the rule requires no such input, but it only licences inferences when the predicates are in fact non-grue? The latter sounds more plausible; otherwise no one who lacked the concept of a gruesome predicate could rationally make inductive inferences. But it raises tricky questions about using that rule in suppositional reasoning. What if (assuming this is possible) different predicates would be gruesome if  $p$  were true, and we are inferring under the supposition that  $p$ ? It isn't obvious just what restrictions should be put on the use of this rule in a suppositional context. Perhaps those restrictions are quite tight. This has important consequences for arguments for the contingent a priori that assume that we can make inductive inferences in suppositional contexts, such as John Hawthorne's *Deeply Contingent A Priori Knowledge* and my *Scepticism, Rationalism and Externalism*.

That was all about the nature of rules; we might wonder whether there is anything that we can say about which rules there are. Different ways of thinking about Carroll's example suggest two different constraints on rules, one more liberal and the other more conservative.

The more liberal constraint is a kind of transcendental consideration. In many cases it seems, at least *prima facie*, that we can get knowledge from a certain source, but we couldn't antecedently justify the use of that source. A classic example of this kind of reasoning is C.A.J. Coady's arguments for against reductionism in testimony. The thought there is that so much of the time the only way we have of checking one person's testimony is through the testimony of another, that if we weren't able to take some testimony as basic knowledge, we'd be led to a debilitating scepticism. That seems unacceptable, so we might take such a rule as given.

The more conservative constraint takes more seriously the particular way in which we need a rule to sidestep Carroll's tortoise. The core problem isn't just that when we add another premise, one that justifies a particular use of the rule, we need yet more to get to the conclusion. Rather, the core problem is that when we add the kind of premise that could justify the rule, we need another step of the very same rule. Justifying this particular use of the rule doesn't seem to get us any closer to where we need to be. Perhaps those cases, where justifying a particular use of the rule still requires the rule, are the only cases where there are unjustified rules.

This seems to be a more conservative principle because whenever it obtains, we'll be able to give a transcendental argument for the existence of a rule. But the converse doesn't seem to hold. We might justify taking someone's testimony about being true because we believe them to be generally reliable. We still need a rule saying that it's good to believe those who are generally reliable, but that doesn't seem like the same rule. Similarly, we might accept (A) as a derived, and hence justified, rule because the best explanation of our experiences is that they are generally reliable. (Jonathan Vogel has argued for this at length over the years.) We still need a rule saying that we should believe the best explanation of a phenomena, but that doesn't look like rule (A) again. On the other hand, Hume's arguments about induction arguably do show that a justification of induction will need to use induction. And a justification of (ND) will, I imagine, still use something equivalent to it. So even on the more conservative conception of rules, they may still be rules.

These are enormous questions, to say the least, so I don't think this goes close to settling anything. But I hope it's enough to suggest that we haven't given up on systematic epistemology once we admit the notion of unjustified rules that justify inferences. And admitting that, which Carroll's example suggests we must, is enough to sidestep

Williamson's argument against EN.

## 5 Science and Evidence Neutrality

If Evidence Neutrality (EN) is true, it is presumably true everywhere. One way to argue against it then is to argue that it doesn't hold in other subjects. And that's what Williamson does. He argues that it doesn't hold in particular in science.

If Evidence Neutrality psychologises evidence in philosophy, it psychologises evidence in the natural sciences too. But it is fanciful to regard evidence in the natural sciences as consisting of psychological facts rather than, for example, facts about the results of experiments and measurements. When scientists state their evidence in their publications, they state mainly non-psychological facts (unless they are psychologists); are they not best placed to know what their evidence is?

If this were a true description of the position of evidence in science, it would be a problem for EN. But it isn't. EN doesn't psychologise evidence in science, it institutionalises it. Let's recall the original statement of EN.

[W]hether a proposition constitutes evidence is in principle uncontentionally decidable, in the sense that a **community of inquirers** can always in principle achieve common knowledge as to whether any given proposition constitutes evidence for the inquiry. (Emphasis added)

Here's one way to preserve EN in a field. Adopt some standards for something being evidence in that field, standards that are in practice (if not always in theory) decidable. Then take questions about whether those standards are good standards to belong to another field. That is, take it that people who are questioning the standards, questioning whether these standards genuinely generate evidence, to be outside the community in the sense relevant to EN. They might of course be part of another intellectual community, but they aren't part of this community. That way we can preserve EN within every given community.

Compare a principle we might call Foul Neutrality (FN) governing a sport. It's pretty important for playing football that we have a quick method for deciding what's a foul and what isn't. And this must be decidable independent of one's interest in the game. We don't get FN by psychologising fouls; we get it by having referees. The referees could be wrong, and indeed we could have interesting projects about improving the quality of referees. But when we engage in that project we've stopped playing football. The community of footballers (as such) satisfies FN because it's part of being in that community that we take the referee's word as final.

Science isn't like football in that it requires absolute respect of the referees judgment. But it is frequently true that the project of using methods or devices to produce evidence is quite distinct from the project of evaluating whether those methods or devices are good. And we can sensibly individuate communities by looking at which methods they take as given. The short version of my response to this argument is that that's really how science works; i.e. that science consists of communities so individuated. Each community has a refereeing institution. Or, at least, it is how it works in the vast majority of cases. In cases where the refereeing institutions break down, where there isn't some other community to serve in effect as referee for your community, then we might have to fall back on psychological states. But EN doesn't systematically psychologise evidence in science.

We might think that evidence must consist of facts measured rather than something about their measurement, because those are the kinds of things we can submit to statistical testing. But that argument, if it works, proves too

much. Williamson's initial description of scientific evidence was that it consisted of the results of experiments and measurements. But that's ambiguous between two readings. On the first, scientists just state the outcomes of their measurements. That is the kind of thing that you can do statistical analysis on. On the second, they state the results of the measurement, and describe what kind of measurement it is. And that's, I think, the true reading. At least for results of any interest, you have to describe how you got them, as well as what you got. But you can't do statistical analysis on a description of a kind of measurement. So it isn't true that all scientific evidence consists of things you can plug into mathematical equations.

On the other hand, this picture of scientific practice does seem to support the institutional picture of evidence. Why is it that we report the methods as well as the result? One simple answer is that it is settled (relative to the kind of science we're engaged with) that using that method produces scientific evidence. That's not to say that the method is beyond dispute. It might be that some other science studies the workings of the very machines that a particular science takes for granted in their operations. It's merely to say that this science has approved the method in question.

We can see this even more clearly if we look at engineering settings rather than science settings. Imagine we're working on a bridge construction project, and we need to know the density of some concrete. We've got a machine that measures concrete density, so we use it and, assuming the answers are plausible, we'll take those answers as given. Evidence Neutrality will be satisfied because we'll agree to use the machine. Of course, the only reason we trust the machine is that there is someone, typically someone else, whose job it is to test the machine on a regular basis, and service it, or have it serviced, if it isn't, and although we might not know the details of how this process works, we'll have a nice certificate saying the machine is in good condition to use. Now the folks who calibrate machines like this aren't perfect, so there are other people whose job it is to audit them on a regular basis. And auditors aren't perfect either, so there will be some body, perhaps a certification body, that oversees them. A positive mark from an auditor only licences a calibrator to approve a machine if the auditor is in turn certified. The board itself may need to be checked, so maybe it will have a board, perhaps including representatives of people like bridge builders who use the machines that we're all interested in.

The crucial point about this story is that at every stage in the process, EN is satisfied. It is similar, I think, in sciences, though the structure is more fluid. Just which sciences will validate the use of the measurement techniques in other sciences is not as straightforward as in engineering. And the precise boundary between questions that are internal to a given science and questions external to it will change over time. When many questions central to the science start to turn on a particular kind of question about measurement, then those measurement questions may become part of the science. (For instance, if experimental philosophy really takes off, perhaps questions about survey design will be regarded as philosophical questions in the future. More prominently, in recent years questions about the behaviour of satellites have become part of climate science because of the importance of satellites to climate measurement.) But still the broad structure is fairly similar.

The big difference between science and engineering is what happens at the end of the process. The way I described the bridge building case was that eventually, the people responsible for checking the activities of others were the very people (or at least the representatives of them) who were being watched over to start with. That obviously isn't what happens in science. We don't check the activities of (say) particle physicists by putting together a board of psychiatrists, nutritionists, economists etc. How might we satisfy EN in basic physics?

Two obvious answers spring to mind. One, either common sense or philosophy tells us that we can take perceptual evidence as given. So even in fundamental physics we can individuate the community in such a way that those who

are raising sceptical doubts are doing something else, namely philosophy.

The other answer is that we might take scientific evidence, at the most fundamental level, to be psychological states. Certainly it isn't uncommon for philosophers of physics to take the role of physical theory to explain our observations. That's part of why we've ended up with such psychologically flavoured interpretations of quantum mechanics, from the Copenhagen interpretation to the many minds interpretation. Perhaps that's just philosophers bringing in bad philosophical prejudices, but it seems like we can do science respecting EN. That's because EN mostly is satisfied by the institutional structure of science, and when it isn't, it doesn't seem to destroy science to take some evidence to be psychological. So there isn't an argument from science against EN.

## **6 Evidence Neutrality as Regulative Ideal**

There is one other argument that Williamson deploys against Evidence Neutrality: it is unattainable. EN requires that the community be able to decide what its evidence is. But an individual can't, in all cases, even decide what her own evidence is. In hard cases, EN doesn't just fail as a theory of group evidence, it fails as a theory of individual evidence.

This isn't something special about evidence. Williamson thinks there is almost nothing that we can, in all cases, tell whether it obtains. Evidence is undecidable because, he argues, practically everything is undecidable in hard cases. The latter conclusion has constraints for norms. If there are norms, then they can't be things that we know to obtain. Williamson gives a nice example. When one is speaking to a group, the rule Adjust the volume of your voice to the size of the room is a good rule, an ideal to aim for, even if we don't know, and can't in principle know, the exact size of the room. Such a norm is a regulative ideal; we aim for it, even if we can't always tell how close we are to hitting it.

So there can be norms that we can't always obtain, or perhaps can at best obtain by luck. EN might, for all Williamson has said, have such a position. We should use evidence that all the members of our community recognise as evidence. The benefits of such a rule can be seen by looking at the relative success, over the course of human history, of individual and group research projects. The great majority of our knowledge of the world is the outcome of research by large, and often widely dispersed, communities of researchers. Even in cases where a great individual advances knowledge, such as Darwin in his theorising about evolution, the individual's work is typically improved by holding themselves to EN as a norm. In Darwin's case, the reason for this is relatively clear, and I think instructive. Darwin collected so much evidence over such a long period of time, that the only way his younger self could convince his later self that it was all part of his evidence was by the same methods that his younger self could convince the community of biologists that it was part of his evidence. It was holding to EN that allowed him to engage in a fruitful long-term research project.

In many ways, EN is quite a weak norm. In earlier posts I discussed what amount to two major exceptions to it. First, EN doesn't require rule neutrality. So the maverick scientist can hold EN while coming to quite bizarre conclusions by adopting various odd rules. As we saw above, we can put some constraints on what makes a good rule, but those constraints won't individuate the good rules. Second, EN, as I'm interpreting it, allows one to choose one's own community. One of the ways we uphold EN in science is by excluding from the community those who doubt the relevant evidence collecting methods. That means we exclude the odd crank and sceptic, but it also means we exclude, from this particular community for the while, those scientists who carefully study the evidence collection methods that we use. In the latter case at least, there is a very real risk that our community's work will be wasted because we are using bad methods. But the alternative, waiting until there is a rigorous defence of a method before we start using it,

threatens a collapse into Cartesian scepticism.

Even if EN is a norm of evidence, a regulative ideal, rather than a constitutive principle of evidence, we might still be pushed hard towards taking intuitions to be evidence. Or at least we might be so pushed some of the time. It doesn't violate EN to take what nutritionists tell us about a healthy diet at face value; the reports of nutrition science are common ground among the community of ethicists. But we can hardly take facts about disputed examples, for instance, as given, even if they are quite intuitive to some of us. And even if, as it turns out, we know the answer. If there are people who are, by any decent standard, part of our community of philosophers, who disagree about the cases, we should be able to give our grounds for disagreement. Not because this is necessary for knowledge, but because the policy of subjecting our evidence to the community's judgment is a better policy than any known alternative.

To be sure, some work needs to be done to show that that taking intuitions as basic does conform to this idea. As Williamson notes, one thing that might (even in somewhat realistic cases) be in dispute is the strength of an intuition. So taking EN as normative might require some modification to intuition-driven philosophical practice. But I don't think it will require as big a diversion as Williamson's preferred anti-psychologistic approach.