Scepticism, defeasible evidence and entitlement

Daniele Sgaravatti

Philosophical Studies
An International Journal for Philosophy in the Analytic Tradition

ISSN 0031-8116
Volume 168
Number 2

Philos Stud (2014) 168:439-455
DOI 10.1007/s11098-013-0134-6

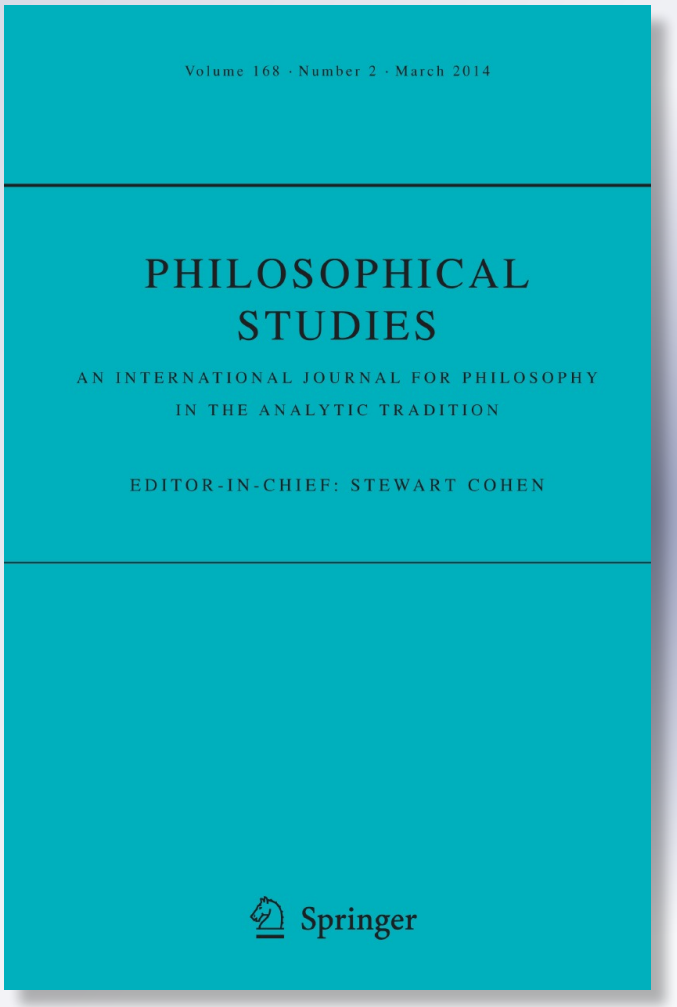

Springer

Philos Stud (2014) 168:439–455
DOI 10.1007/s11098-013-0134-6

Scepticism, defeasible evidence and entitlement

Daniele Sgaravatti

Published online: 26 April 2013

Abstract The paper starts by describing and clarifying what Williamson calls the *consequence fallacy*. I show two ways in which one might commit the fallacy. The first, which is rather trivial, involves overlooking background information; the second way, which is the more philosophically interesting, involves overlooking prior probabilities. In the following section, I describe a powerful form of sceptical argument, which is the main topic of the paper, elaborating on previous work by Huemer. The argument attempts to show the impossibility of defeasible justification, justification based on evidence which does not entail the (allegedly) justified proposition or belief. I then discuss the relation between the consequence fallacy, or some similar enough reasoning, and that form of argument. I argue that one can resist that form of sceptical argument if one gives up the idea that a belief cannot be justified unless it is supported by the totality of the evidence available to the subject—a principle entailed by many prominent epistemological views, most clearly by epistemological evidentialism. The justification, in the relevant cases, should instead derive solely from the prior probability of the proposition. A justification of this sort, that does not rely on evidence, would amount to a form of *entitlement*, in (something like) Crispin Wright's sense. I conclude with some discussion of how to understand prior probabilities, and how to develop the notion of entitlement in an externalist epistemological framework.

Keywords Scepticism · Consequence fallacy · Entitlement · Tim Williamson · Crispin Wright

D. Sgaravatti (✉)
Università Dell'Aquila, Viale Nizza 14, 67100 L'Aquila, Italy
e-mail: daniele_sgaravatti@yahoo.it

1 Introduction

My main aim in this paper is to assess a certain very general form of sceptical argument, and to explore a certain sort of response to it, one involving giving up an initially plausible link between evidence and justification. In the second section, I approach the issue from a slightly unusual angle, describing a certain sort of fallacy, first discussed by Williamson. The fallacy consists of a certain kind of misguided reasoning which leads to the rejection of any sort of non-deductive argument. I attempt to clarify what the fallacy consists in (correcting, as I see it, some criticisms expressed recently which are based on a failure to interpret correctly the fallacy), and I then describe two ways in which one might commit the fallacy. The first way, which is rather trivial, involves overlooking background information; the second way, which is the more philosophically interesting, involves overlooking prior probabilities. In the third section, I will introduce a general form of sceptical argument, building on work by Huemer, against the possibility of defeasible justification. That argument would be in effect a vindication of the sort of reasoning we called the "consequence fallacy". I will individuate three general principles that the argument depends upon. One of the principles, which was missing in Huemer's presentation of the argument, says, roughly, that a belief cannot be justified unless it is supported by the totality of the evidence available to the subject. This principle is entailed by many prominent epistemological views, most clearly by epistemological evidentialism. Nevertheless, I argue that there are good reasons to give up that principle in order to resist the sceptical argument. The justification of the belief, in the cases for which the principle fails, should derive not from the evidence but rather from the prior probability of the proposition. In other words, I will argue that we can resist the sceptical argument, provided that prior probabilities yield a form of warrant that does not rely on evidence, so a form of *entitlement*, in (something like) Crispin Wright's sense. In the fourth section, I will discuss whether there is an understanding of prior probabilities that makes them suitable to play this crucial role, and I will conclude with a tentative discussion of how to develop the notion or prior probability and the notion of entitlement in an externalist, and even "knowledge-first", epistemological framework.

2 The consequence fallacy

Williamson (2007) describes what he calls the 'consequence fallacy'. He claims that such a fallacy plays a large role in some sceptical arguments, and in particular that the failure to recognize it as a fallacy could lead to a form of scepticism about inductive support. As far as I know, Williamson's discussion attracted very little attention from epistemologists,[1] perhaps because it occurs in a book which is not mainly about epistemology. But I think the discussion of this problem will prove interesting in itself, and it will illuminate a certain important form of sceptical puzzle.

[1] The only notable exception that I am aware of is Wong et al. (2012).

In this section, I will describe the "consequence fallacy", and I will try to understand it properly. The notion is introduced in Williamson (2007), using a probabilistic framework, by the following passage:

> Let e be a body of evidence which raises the probability of a hypothesis h to a value close to 1 without quite making h certain, so $Pr(h) < Pr(h \mid e) < 1$. The material conditional $e \supset h$ is a logical consequence of h, and therefore at least as probable as h; in fact, $Pr(e \supset h \mid e) = Pr(h \mid e)$. However, e is evidence against $(e \supset h)$, for $Pr(e \supset h) > Pr(e \supset h \mid e)$, simply because $e \supset h$ is true in all those possibilities which e eliminates ($e \supset h$ is a logical consequence of $\sim e$). Clearly, all of this is compatible with a high degree of legitimate confidence in both h and $e \supset h$. Whenever evidence makes some hypothesis more probable than before without making it certain, that evidence makes some logical consequence of that hypothesis less probable than before. (···) What this reveals is a fallacy in the tactic of criticizing confidence in a theory by identifying a logical consequence of the theory (not itself a logical truth) whose probability is not raised by the evidence. Call that the *consequence fallacy*. (Williamson 2007 p. 232, adapted notation).

I'll take the latter lines to provide the 'official' definition of the fallacy. Some clarification is however needed. Wong et al. (2012) state two different readings that could be given of the argumentative strategy Williamson is describing there. They are the following:

> (T1) Arguing that the probability of a theory is not raised by the evidence by identifying a logical consequence of the theory whose probability is not raised by the evidence.
>
> (T2) Arguing that a theory is not made likely to be true by the evidence by identifying a logical consequence of the theory that is not made likely to be true by the evidence (Wong et al. 2012 pp. 1–2)[2].

(T2) is different from (T1) because a hypothesis can have a high probability on some evidence (which is all Wong and Yudell mean, somewhat misleadingly, by saying that the evidence makes the hypothesis likely to be true) even though the evidence does not raise the probability of the hypothesis. This point will be crucial for me as well in the interpretation of the fallacy. Wong and Yudell go on to argue that while the argumentative strategy described in (T1) is indeed a fallacy, it is not a commonly used strategy. On the other hand, they claim that the argumentative strategy described by (T2), while it is sometimes used by philosophers, is not a fallacy. I think they are right on all counts about (T1) and (T2). But neither (T1) nor (T2) is a plausible reading of what Williamson writes. Remember what (T1) and (T2) are trying to render is "criticizing confidence in a theory by identifying a logical consequence of the theory (···) whose probability is not raised by the evidence". Wong and Yudell write: "Williamson's use of the phrase 'probability is not raised' suggests (T1), while his use of the word 'confidence' suggests

[2] Page numbers refer to the online manuscript.

(T2)."(Wong et al. 2012, p. 2). I think they are right on each expression; the problem is with their implicit assumption that the two expressions must refer to the same thing. We should simply take each expression to mean what it seems to mean. The result will combine (T1) and (T2) as follows (adopting Wong and Yudell's reading of likely to be true):

(CF): Arguing that a theory is not made likely to be true by the evidence by identifying a logical consequence of the theory whose probability is not raised by the evidence.

This was of course very abstract. While my focus will be on philosophical applications of the consequence fallacy, I think something similar enough can occur in everyday contexts as well, where of course we will not find explicit talk of probabilities. A couple of examples of the latter sort will be useful. Suppose I am told by someone that her name is Anna. Should I believe that her name is Anna? I reason as follows; if her name is really Anna, then she was not lying about her name being Anna. But surely, from the fact that she said that her name was Anna, I cannot conclude that she was not lying about her name being Anna. So I refrain from forming any belief as to her name. In a second case, I see that my fuel gauge points to empty. I am tempted to conclude that I am out of gas. However, I reflect, I could then conclude that, in this occasion, it is not the case that the gauge points to empty but the tank is not empty, on the basis of the fact that the gauge points to empty. But surely I cannot do that; whatever the gauge indicates is not by itself evidence of its reliability. So I refrain from forming the belief that I am out of gas.

In both cases, there is evidence for a seemingly innocent hypothesis (the person's name is Anna; the tank is empty), but the hypothesis has a consequence that is not supported by the evidence (the person did not lie about being called Anna; the gauge is not malfunctioning). Why is it a fallacy to reject the hypothesis and its consequence? For one thing, I might know independently that the suspicious consequence holds. Let us focus on the second case. I might know that the gauge is reliable independently of what it points to in this occasion. I might have had experience of its reliability, or of the reliability of fuel gauges of that kind, or, even more generally, of the reliability of indicators built by humans. Of course, I might know instead that this gauge is often unreliable, which would make my reasoning less peculiar. This just comes down to the familiar point that what a certain proposition, e.g. that the gauge points to empty, is evidence for, will depend on my background information. In more abstract terms, I might be able to infer h from e, while I have background evidence for $e \supset h$.

So this is a way in which one might commit the consequence fallacy. This reconstruction is attractive because it explains, at the same time, why the reasoning is flawed, and why we might find it attractive. The answer is the same: because we are forgetting the role of background information.

It would be hasty however to conclude that this is all there is to the consequence fallacy. This is not what Williamson had in mind. Of course, he would not deny that you can forget about background information. Still, he would insist that we can include all the background information in the evidence, and still get the same fallacy. For suppose our evidence e, so constructed, supports, although it does not

entail, a proposition h; there will then be a logical consequence of h, namely e ⊃ h, the negation of which entails all of our evidence. Still, we could be justified in having a high degree of confidence in both h and e ⊃ h.

So, one would like to ask, how are we entitled to a high degree of confidence in e ⊃ h? Williamson does have an answer to that question, albeit he is not crystal clear about it. I think it's something like the view expressed in this passage:

"On some views, if the prior probability of p is high enough, we should be confident of p even if its probability is somewhat lowered by Ap [the appearance of p]" (Williamson 2007 p. 228).

The crucial notion is *prior probability*. Although I am presupposing some familiarity with the probability calculus, a few words are in order here on the notion. The prior probability of a proposition is the probability that is assigned to that proposition before the acquisition of some relevant evidence. Once the probability is updated on the evidence, the result may be considered as a new prior probability with respect to the acquisition of further evidence. Therefore the prior probability will usually incorporate the influence of some background evidence. However, it is interesting, philosophically, to ask what the prior probability of a proposition should be before the acquisition of any evidence. Henceforth, 'prior' will refer to the probability a subject gives or should give to a proposition prior to, and independently of, all of the subject's evidence. I will come back to this in the last section.

As I said, Williamson is not explicit about embracing the explanation of the fallacy in terms of prior probabilities. However, it seems to me the only one, excluding the background information, which is compatible with a probability function which respects the standard set of axioms for probability calculus, as Williamson explicitly requires[3] of the sort of probability he is talking about.[4] As Williamson notes, $Pr(e \supset h \mid e) = Pr(h \mid e)$, and $Pr(e \supset h) > Pr(e \supset h \mid e)$. But, if we are entitled to a high degree of confidence in h on e, and the prior probability of e ⊃ h is higher than the probability of h on e, clearly we are entitled to a high degree of confidence in e ⊃ h on the basis of its prior probability.

So here is another way we could explain the consequence fallacy. Again, we have a single answer to the questions 'why is the reasoning in the consequence fallacy flawed?' and 'why is it an attractive error?'; the answer is now: because we are forgetting the role of prior probabilities.

Of course the two explanations are in many ways compatible. First, different cases might be better explained by the one or the other. Secondly, we might forget in a single case about both background information and prior probabilities. Finally, there might be ineliminable relations between what evidence a subject possess and what prior probabilities the subject does or should adopt. However, because, as I said, we will consider cases in which the totality of a subject's evidence is taken into

[3] Williamson (2007), p. 228, refers back to Williamson (2000), and the latter, p. 211, makes explicit the assumption of standard (Kolmogorov) axioms for probability calculus. See also n.11 for my own adherence to the standard calculus.

[4] We will come back to the nature of this sort of probability.

account, background evidence will not be able to play any direct role in what follows.

In the next section I will discuss a sceptical argument against the possibility of defeasible justification, one that, if successful, would indeed show, among other things, that the "consequence fallacy" is not fallacious. I will argue that appreciating the role of prior probabilities can help us to resist that argument.

3 Consequence fallacy and a Humean puzzle

I agree with Williamson that proper understanding of the consequence fallacy has a role in evaluating many sceptical arguments, although I disagree with him on what lesson we should learn from recognition of the fallacy. I will consider just one elaboration of a sceptical argument[5] , which seems particularly relevant, both in itself and in this context. It is an argument presented, although not endorsed, in Huemer (2001), for the impossibility of defeasible justification. Like Huemer, I will take 'defeasible justification' to mean justification given by evidence that does not logically entail the (allegedly) justified proposition or belief. This argument constitutes, I think, a Humean sceptical puzzle. I would go as far as saying that it captures the essence of Humean scepticism, targeting the very idea of defeasible justification.

Huemer limits the argument to 'inferential' justification; however, all he means by that, as far as I can see, is justification grounded in propositional evidence. If one thinks, as Williamson does, that all evidence is propositional, one can put aside this limitation and generalize the argument. Even if one does not believe all evidence to be propositional, however, I think the argument could be reformulated in such a way to apply to non-propositional evidence as well. A suitable notion of 'entailment' however would have to be developed, presumably in modal terms, so that the evidence, whatever it consists of, would "entail" a proposition for a subject when it is not possible for the subject to have that evidence and the proposition to be false. For simplicity's sake, I will stick to the assumption that all evidence is propositional, and I will reformulate Huemer's argument accordingly. It is also important to note that the argument proceeds in terms of the notion of justification (in particular, it has to be understood as about propositional[6] justification), while Williamson's discussion was about entitlement to a certain degree of confidence. I will come back to this difference later.

Say we have defeasible justification when the set of propositions constituting your total evidence, call it e, supports the conclusion, call it p, although e & $\sim$p is logically consistent. The support given by our evidence should be at least strong enough for the subject to be justified in believing p; and strong enough, in absence of other obstacles, for that belief to qualify as knowledge.

[5] Weatherson (2007) and Schiffer (2004) also present forms of sceptical argument for which the following considerations are relevant.

[6] In the sense of Firth (1978).

Huemer's argument to the effect that this is never going to be the case uses two principles. I will re-formulate the principles rather radically, with respect to Huemer's version, eliminating (what seem to me) some unnecessary complications,[7] as follows:

P1 Given any evidence, *e*, and hypothesis *h*, if *e* is justification to believe *h* for a subject *s*, *s* must have, given *s*'s total evidence, justification to believe each logical consequence of *h*.

P2 For any sets of propositions *e* and *h*, if *h* entails *e*, then *e* is not a justification to believe the negation of *h*.[8]

Here is an informal statement of the argument. Consider a particular case in which my evidence is e = 'I have perceptual experience as of a hand in front of me', and the desired conclusion is the hypothesis h = 'I have a hand'. The sceptical hypothesis is simply sh = 'I have perceptual experience as of a hand in front of me, but it's not the case that I have a hand'. Since sh entails e, by P2 it is not the case that e is a justification to believe not-sh; since sh is incompatible with h, by P1 e cannot justify h, unless we have a justification to believe not-sh in some other way. Do we have different evidence for rejecting sh? If we do, let's add it to our initial evidence e, and get the new evidence e*. If e* rules out every possibility incompatible with h, then we have an indefeasible justification after all. Suppose however e* still does not entail h. We can then consider the new sceptical hypothesis sh* = 'e* but it is not the case that I have a hand'; and of course we can apply the same reasoning. There seems to be no stop to this regress. In a more general form:

Suppose e is your total propositional evidence.

1) Your total evidence e does not entail h (assumption)
2) There is an hypothesis h* (= e & not-h) whose negation is a consequence of h, and which entails e (from 1)
3) You do not have a justification to believe the negation of h* (from 1), 2) and P2)
4) e is not a justification to believe h for you (from 2), 3) and P1)

In fact, the argument as it stands is not valid. The problematic passage is from (2) to (3). By P2, since h* entails e, e is not a justification to believe the negation of h*. But, by (1), e is our total evidence; still, the subject might have a justification to reject h* that is not part of the subject's total evidence (and cannot be added to it, as the informal presentation of the argument presupposed instead, perhaps because it is not of propositional nature). In other words, the argument goes through only if the following presupposition, which I think is therefore worth considering a third principle, is in place:

P3. For any subject *s* and hypothesis *h*, if *s*'s total evidence is not a justification to believe *h*, than *s* has no justification to believe *h*.

[7] In particular, Huemer employs the notion of *reason* to believe, without making clear in what it differs, if it does, from the notion of justification to believe.

[8] P2 entails, assuming classical logic, that where h is a logical contradiction we cannot have reasons to reject it. If this seems wrong, we can consider P2 as implicitly limited to exclude such cases. I tend to think that P2 is fine as it is, and we do not need reasons to reject contradictions.

Like P1 and P2, P3 seems initially very plausible. It is also a widely held view. It is arguably entailed for example by epistemological evidentialism. Evidentialism, as described (and defended) by Feldman and Conee (1985), is the following view:

"Doxastic attitude *D* toward proposition *p* is epistemically justified for *S* at *t* if and only if having D toward *p* fits the evidence *S* has at *t*" (Feldman and Conee 1985 p. 15). P3 is more restricted than evidentialism, and it is only concerned with one direction of the bi-conditional linking evidential support and justification; it says that your total evidence constituting a justification is necessary for you to have a justification, but it does not say that it is sufficient. Still, if P3 is entailed by evidentialism, one would need to give up the latter if one gives up the former. It should also be noted that evidentialism, as we have defined it here, is a very general view, and it is not for example tied to epistemological internalism of the sort defended by Conee and Feldman, or of any other sort. P3 is also plausibly entailed by the evidentialist reliabilism defended by Comesana (2010), and it could be endorsed by someone who has an externalist conception of evidence, like the one defended by Williamson[9] or the one tentatively proposed in Goldman (2009). Even more generally, it might be thought that, given that we can allow this extremely broad reading of "evidence", P3 is entailed by any form of epistemic foundationalism. Nevertheless, I will end up suggesting that we should reject P3.

Huemer's argument (or my version of it) is connected to the consequence fallacy as follows (again putting aside the distinction between justification and support): it moves from the logical possibility of e & $\sim$h, which is of course equivalent to the logical possibility of the negation of the material conditional e $\supset$ h, to the conclusion that e does not support h. While Williamson noted that endorsing the argument form he dubs the "consequence fallacy" would have led to the conclusion that no defeasible justification is possible, the argument Huemer presents seems to provide an independent reason to think that the argument form is indeed not flawed, and it concludes that no defeasible justification is possible.

If one accepts P1 to P3, one has to deny that the consequence fallacy is a fallacy, and to embrace the sceptical conclusion that defeasible justification is impossible. This would arguably be a very radical[10] sceptical conclusion. Let us therefore look at some ways of rejecting the argument. We face of course a choice, since we have to reject at least one among P1 to P3. I will discuss the three options in turn, and I will argue that P2 cannot be abandoned, and also that a weaker version of P1, which seems not negotiable, is sufficient for the argument to go through; therefore we should reject P3, unless we are willing to embrace the sceptical conclusion.

P1 seems tantamount to a principle of closure for propositional justification. I will put it aside for a moment, supposing it has enough initial plausibility. If we assume

[9] It is somewhat hard to tell whether Williamson would actually endorse P3, since P3 uses the notion of justification, which is not part of Williamson's official ideology. In Sect. 4 of the paper I give some reasons to think however that he *shouldn't* accept P3.

[10] How radical depends on how much indefeasible evidence we typically possess. Pre-theoretically, there is a sense in which much of our scientific and everyday beliefs are based on defeasible evidence. I cannot discuss the issue here. However, even someone like Williamson, who thinks that indefeasible evidence is abundant (though see Dodd 2007), needs to deny one of P1–P3; or so I argue at the end of this section.

P1, we have to individuate the principle to be rejected among P2 and P3. However, denial of P2 could be defended only on a peculiar reading of 'being a justification to believe'. If h entails e, then e cannot make h less likely than it was[11]; if so, it might still be the case that one has a justification to reject h, but it does not seem at all plausible to say that e itself constitutes such a justification. Rather, in that situation one has a justification to reject h *despite* the support offered by the evidence for h.

Williamson, if I read him correctly, would opt instead for rejecting P2. This is related to his discussion of what he calls 'appearance principles'. An appearance principle is something similar to the following: "one should be confident that p (on the basis of common sense) only if its appearing (to common sense) that p is good evidence that p" (Williamson 2007 p. 227). Different appearance principles can be obtained by substituting common sense with other sources of belief (such as perception). Williamson claims that appearance principles lead to scepticism. Consider the sceptical hypothesis relative to some proposition p that an evil demon makes it the case that all appearances indicate that p is the case, but p is false. Call that hypothesis SS*. Williamson notes that the appearance that SS* does not hold is not truth-indicative. Therefore, he claims, "by the relevant appearance principle, one should not be confident that SS* does not obtain" (Williamson 2007, p. 230). But no appearance principle of the form described above delivers the general, unrelativized, result that one should be confident that SS* does not obtain. They will only deliver the result that you should not be so confident on a certain basis. Williamson might be thinking that, given that we are in fact confident that SS* does not obtain, we are bound to be so confident on the basis of the totality of the appearances available to us, or the totality of our evidence (note that this is implicitly assuming P3), although that totality does not raise, and in fact lowers, the probability that SS* does not obtain. This would then be a counterexample to P2; although SS* implies all of our evidence, we are justified to believe its negation on the basis of that evidence. However, this conclusion seems to be acceptable only on an extremely unnatural reading of 'on the basis of', related to the unnatural reading of 'being a justification' noted above. To see that this is not our usual understanding of 'on the basis of', consider the following analogy. Suppose I am told to lower the volume of my TV, and I then proceed instead to raise it. I then claim to have acted *on the basis of* what I was told to do. When someone objects, I reply as follows: 'Admittedly, I already had a strong inclination to raise the volume; therefore, I would have done that on the basis of a large class of different events. Not anything would have done though; if I had been told to help myself to a piece of cake, for example, I would have done that instead. Therefore, I raised the volume of the TV on the basis of the fact that I was told to lower it.' I don't think anyone should be convinced by my reply. I acted on the

[11] I am taking this to be a strong pre-theoretic judgement; if the evidence is exactly what the hypothesis entailed is would be, the hypothesis has done as well as it could. So I do not regard that assumption as the product of a specific formalization of probability calculus; rather, the other way round, standard probability calculus is plausible (in this respect at least) as a way of representing evidential probability because it respects this pre-theoretical commitment. If one does not share this commitment, one can surely find a way of adapting probability calculus in order to be allowed to flout it. See Weatherson (2007), for a technically and philosophically ingenious attempt.

basis of what I was inclined to do, and not of what I was told. But I don't see any relevant difference in the epistemic case. In a situation in which all of my evidence lowers the probability of a certain hypothesis, I might still rationally give it high credence; but it makes no sense to say that my evidence is the basis of my doing so. My evidence would be a reason or a justification to do the opposite, if anything. The negation of P2 therefore still looks very implausible.

However, if accepting P2 inevitably led to the consequence that defeasible justification is impossible, thus making the consequence fallacy not a fallacy at all, that would itself be a strong reason against P2. But we have a better alternative, which is denying P3. Here is a first rationale for this choice. P3 says that if the total evidence of a subject does not constitute a justification, then the subject has no justification to believe something. Because the prior probability is not itself part of the evidence, this implies that prior probabilities cannot provide reasons to reject or accept a hypothesis; but we have seen that whenever a body of evidence supports a hypothesis without entailing it, there is a consequence of the hypothesis which we are entitled to a high degree of confidence about. Because we are already considering the totality of our evidence, this will not be a case in which background evidence is playing a role, but rather one in which prior probability is crucial.

Here is another way of getting to the same conclusion, which will also show why rejecting P1 would not get to the heart of the problem. The work done by P1 in the sceptical argument can also be done by the following probabilistic closure principle for degrees of justification, which seems (unlike P1) rather uncontroversial:

CJ. Given any evidence *e*, and hypothesis *h*, if *e* is justification to believe *h* to a degree n for a subject *s*, *s* must have, given *s*'s total evidence, justification to believe each logical consequence of *h* to a degree m $\geq$ n.

CJ avoids an important class of problems for closure principles, those related to lottery puzzles. Like other closure principles for propositional justification, it requires a certain level of idealization, but apart from that it looks in good standing. It seems very hard to imagine we can have a probabilistic representation of epistemic support without something like that in play. We can reformulate the rest of the argument in such a way that it only requires this weaker closure principle. Here is how:

Suppose T (for threshold) is the level of epistemic probability of a hypothesis necessary for justification. If however you are justified in believing some hypothesis h which is not entailed by your evidence, and hence h has a degree of probability above T, by CJ, the consequence of the hypothesis e $\supset$ h also must have a degree of probability above T. But the level of probability of e $\supset$ h cannot have been raised by your total evidence e, because, as we saw above discussing principle P2, the only cases in which e $\supset$ h fails are cases in which e holds. Therefore e $\supset$ h had a degree of probability above the threshold T independently of the total evidence, on the basis of its original prior probability, which entails the negation of P3. In other words, the negation of P3 follows from recognizing the consequence fallacy as such, and from understanding the role of prior probabilities in it.

As I noted above, Williamson actually never talks of justification to believe, but only of entitlement to (by which I take he simply means justification for) a certain degree of confidence. Williamson in a way might be willing to accept, and indeed

might be committed to, the conclusion that there is no defeasible justification,[12] insofar as he holds that one should believe only what one knows, and what one knows on his view has maximum epistemic probability. However, he certainly wants to avoid the result that there is no defeasible support; that is, the conclusion that you are not entitled to any degree of confidence in anything which is not entailed by your evidence. We will be able to avoid this result, however, only if some hypothesis h has a high probability n between 0 and 1 on the totality of your evidence e, and the material conditional $e \supset h$ has a prior probability equal to or higher than n.

It is now time to give a further look to the notion of prior probability and its relation with the notion of entitlement.

4 Prior probabilities, entitlements and the sceptical challenge

Williamson (2000), in describing the kind of probability that is also used in Williamson (2007), writes: "The discussion will assume an initial probability distribution P. P does not represent actual or hypothetical credences. Rather, P measures something like the intrinsic plausibility of hypotheses prior to investigation; this notion of intrinsic plausibility can vary in extension between contexts." (Williamson 2000, p. 211).

This is almost everything which is said about prior probabilities, except that "the characterization of the prior distribution of evidential probabilities is blatantly vague" (Williamson 2000, p. 212).

This of course is not incredibly encouraging. However, it is worth noting Williamson's remark on the variation between "contexts". I am not at all sure what 'contexts' are here. My supposition would be that Williamson is not thinking of conversational contexts in the sense in which some epistemologists claim that the meaning of 'to know' varies among contexts. Rather, I think he has in mind that there will be a different intrinsic plausibility, and so different probability distributions will be reasonable, with respect to different sets of evidence. This would perhaps make sense on one's being entitled to a high degree of confidence (or one's having reasons to reject) on the basis of some evidence, when the entitlement seems to derive from the prior probability. Still, it is very unclear how this could work. I will come back to the suggestion however at the end of this section.

Something more needs to be said on the notion of prior probability, since the only interpretation of the consequence fallacy (independently of whether or not one rejects acceptance principles and related principles) which makes sense of radical sceptical arguments, and which respects standard axioms of the probability calculus, requires prior probabilities to play a decisive role in excluding sceptical hypotheses.

In particular, in some circumstances, it must be the case that the totality of our evidence lowers the probability of some hypothesis, but we are anyway entitled to a high degree of confidence about that hypothesis. This seems to me to have the

[12] In the sense of 'defeasible' which I am employing. In a different sense, Williamson explicitly holds that even knowledge is defeasible, in that it can be lost due to one's acquiring new, misleading, evidence (see Williamson 2000 pp. 205–6 and 218–21).

interesting consequence that, given that prior probability, we would be entitled to a high degree of confidence in that hypothesis on no evidence at all, or, if you prefer, on evidence consisting solely of the empty set. The sceptic however might challenge that idea. If intrinsic plausibility is in some sense independent of evidence, it might be argued, it should be a priori; but it is not clear that the sort of proposition involved could be supported a priori, since they typically are, if true, only contingently so.

Williamson anticipates how a sceptic would react to this idea in the continuation of the quote I had above: "On some views, if the prior probability of p is high enough, we should be confident of p even if its probability is somewhat lowered by Ap [the appearance of p]. The judgement sceptic regards such a defence of disputed philosophical propositions as unacceptably dogmatic, having the advantages of theft over honest toil."[13] (Williamson 2007 p 228).

How legitimate is the sceptic's complaint? The complaint cannot be that all atomic contingent propositions should have equal prior probability; that would seem just as arbitrary as any other assignment, and it is not even clear that such a strategy can be consistently applied.[14] I agree with Williamson (2000) on two negative points: prior probabilities cannot be just subjective credences, nor can the probability function work on purely syntactic grounds. I cannot argue for this negative point here. Let me just note that the most prominent alternative, the subjective interpretation, is usually thought to have difficulties to handle the particular problem we are discussing here, the nature of prior probabilities. For the sceptic will require the priors to be justified, if they are to ground justified beliefs. But interpreting priors as simply credences seems to give no guarantee that they are in any sense justified, unless one embraces a generalized form of epistemic conservatism, on which subjective credence in a proposition, provided it is coherent with the rest of the subject's credences, yields a correspondent level of epistemic warrant for the proposition.[15]

Even if the negative point is granted, however, it only makes the situation worse, since no well worked-out alternative notion of prior probability is in sight. Williamson observes that the lack of an operational definition is not a good reason to avoid using a notion. I certainly agree that a lot of significant work can be done just taking the notion for granted. However, in some sense, the complaint seems legitimate; more honest toil is needed on the notions of 'prior probability' and 'intrinsic plausibility', if we are not just to point out the fallacy, and call it such (which is of course a necessary step), but to give an explanation of why it is a fallacy. In particular, what is needed, it seems, is something like Wright's notion of *entitlement* (Wright 2004); a form of warrant which comes "for free", independently of the evidence.

Before looking more closely at this notion, let me clarify the reasoning that led me to claim that the notion is needed. One may try to show that the consequence

[13] A "judgement sceptic" is someone who is a sceptic about human judgements in general.

[14] Some think that the lesson of Bertrand's paradox is precisely that we cannot assign probabilities this way. See for example Shackel (2007).

[15] See Harman (1986) and (2003) for some defense of epistemic conservatism, and Swinburne (2001) for a defense of a qualified version of this view applied to epistemic probability.

fallacy is not fallacious, and hence to defend scepticism, through the following argument:

1) Suppose on the basis of e I acquire a justified belief in h, although e does not entail h.
2) Since e ⊃ h is a consequence of h, and I have a justified belief in h, I must be in a position to infer e ⊃ h from h, thus acquiring justified belief in e ⊃ h.
3) My belief in h is supported by e, so if I acquired a justified belief in e ⊃ h through inferring it from h, my belief in e ⊃ h would also be supported by e.
4) Since e lowers the probability of e ⊃ h, it cannot support e ⊃ h.

1–4 are clearly inconsistent. A sceptic about the possibility of defeasible justification would then conclude that the supposition expressed by 1 has been reduced ad absurdum.

Williamson, whatever he thinks of 2 and 3, would probably reject 4[16]; my diagnosis is that the culprit is 2. But this is just to say that the reasoning is a case of failure of warrant transmission, in Wright's sense. Closure demands that if I have justification to believe h, and h entails e ⊃ h, then I have justification to believe e ⊃ h; it does not demand that I am able to acquire a justification for e ⊃ h through h, because to be justified in believing h might require already being justified in believing e ⊃ h in the first place. This shows that, supposing that it is not impossible to acquire justified belief (as opposed to mere high degree of rational confidence) through defeasible evidence, what the prior probability has to provide for the consequence of the belief is nothing less than justification. Moreover, the proposition involved will typically be a contingent one.

If, like Williamson, we framed the discussion in terms of high degree of confidence, the role of prior probabilities would of course be correspondingly different. On the interpretation I am currently exploring, Williamson's view denies that it is possible to get full justification on the basis of defeasible evidence, thereby making a significant concession to scepticism. However, the view will still encounter difficulties substantially similar to the ones we discussed. The role of prior probabilities in providing a high degree of confidence would still be crucial, and just as mysterious. If my total evidence entitles me to degree of confidence n in a proposition, simple probabilistic reasoning requires that the prior probabilities justify a degree of confidence (at least) equal to n in the material conditional having the total evidence as antecedent and the proposition in question as a consequence. How can we be entitled to a high degree of confidence in a (contingent) proposition on no evidence? This question is not avoided.

I do not have space here to discuss at proper length the notion of entitlement, or a probabilistic version of it. But I would like to close with some considerations on how such a notion could, and should, be developed, which are suggested by the foregoing discussion. I am using 'entitlement', as I noted above, in a very broad way, to cover all forms of positive epistemic status which are independent of the

[16] It seems to me the negation of 4 is entailed by the denial of 'appearance principles'.

evidence possessed by a subject.[17] In this raw form the notion, although not under the name of 'entitlement', has received a certain amount of attention, even independently of Wright's ingenious elaboration and defence, for example in Stine (1976), Cohen (1988), Field (2000), and Harman (2003). However, all of these attempts at developing the notion of entitlement, including Wright's, have an epistemologically internalist character. More specifically, most of the authors I mentioned make the possession of entitlement in some sense depend on the subject: for example, on which epistemic project the subject is engaged in (Wright); on what epistemic values the subject endorses (Field); on which beliefs the subject has (Harman); or on what the subject is presupposing in a context (Stine). Cohen comes somewhat apart where he suggests that the belief in the negation of a sceptical hypothesis possesses an "intrinsic rationality" (Cohen 1988 p. 112). But he does not do much to explain what this feature consists in, and why some beliefs have it and others do not.

If I am right that Williamson needs the notion of a form of justification based on the prior probability, and a fortiori a notion of entitlement (justification independent of the evidence), clearly he would have to develop that notion in a somewhat different way, given the epistemological framework of Williamson (2000). I believe that the best available suggestion is in Ichikawa Jenkins (forthcoming).[18] Ichikawa Jenkins explores, although she does not endorse, the strategy of appealing, in solving some epistemological underdetermination problems (not structurally dissimilar to sceptical puzzles of the sort I considered), to what she calls "justification magnetism" (by analogy to the Lewisian idea of reference magnetism; see e.g. Lewis (1984)), the idea that some propositions *attract* justification because of the sort of content they have, and independently of what a subject's evidence is, or even, perhaps, of our epistemic practices as a whole. The following remarks should be understood in the same exploratory spirit adopted by Ichikawa Jenkins. The idea of justification magnetism seems to fit very well with Williamson's claim, quoted above, that prior probabilities should represent the "intrinsic plausibility" of a proposition. But what intrinsic property of propositions could play the desired role? Ichikawa Jenkins considers several alternative suggestions. One she discusses at some length is that just the fact that a proposition is true could play the desired role. The foregoing discussion seems to provide two worries for this suggestion (to be added to the ones discussed by Ichikawa Jenkins). Ichikawa Jenkins notes that, unlike Lewisian naturalness, truth, standardly conceived, does not come in degrees. However, whatever property determines intrinsic plausibility, it should be a property that comes in degrees, to match the degrees of the prior probabilities. Moreover, if we wish to respect the insight that we can rationally use fallible methods, it would seem desirable to allow that we may sometimes be justified in giving a high degree of confidence to a false proposition. In that situation, the material conditional having our evidence as antecedent and the false proposition as

[17] The notion of 'evidence' however should also be understood in a broad way, so that it encompasses not just all the propositional evidence but also, if there is some, non-propositional evidence provided by experiential states. This differentiates my use of 'entitlement', for example, from the use of Burge (2003).

[18] My discussion is based on the pre-publication version of the manuscript.

consequent will plausibly be false, but we need the prior probability to allow a high degree of confidence in it as well.

Another suggestion Ichikawa Jenkins considers is appealing to a notion of "naturalness". But such as a notion seems too much in need of being spelled out to be more than a label.[19] A further suggestion is that we could exploit at least in part the apparatus of Wright (2004). In Wright's view (or the most prominent version of it), we have entitlement for propositions of a particular kind, those that are, first, presuppositions of a cognitive project; where some proposition is a presupposition of a cognitive project when doubting the truth of that proposition rationally requires doubting that the project is significant and/or feasible. Second, the proposition must meet two further conditions. It must be such that we have no reasons to believe that it is false, and, crucially, such that the attempt to find reasons to accept it would involve a cognitive project with further presuppositions no more certain than the proposition itself. Furthermore, Wright poses some constraints of the sort of cognitive project whose presuppositions can generate entitlement; such project must be "indispensable, or anyway sufficiently valuable to us" (Wright 2004 p. 192).

Now, as Ichikawa Jenkins also notes, if we were to say a proposition "attracts" justification precisely when Wright's conditions are met, this could be seen simply as a terminological variant of Wright's view. But I believe that there is room for an externalist version of the view.[20] I will end by sketching it. The crucial modification should occur in the last element of Wright's picture that I described, the factors determining the relevance of the cognitive project. I believe this could be evaluated externalistically, as a function, for example, of the amount of knowledge that the subject would be able to acquire through that project. This would allow us to provide a graded notion. The degree of confidence in a proposition that we are entitled to would be proportional to the cognitive fruitfulness of the project. Note also that how much knowledge one can derive from a cognitive project will depend very strongly on what knowledge or evidence (which is the same for Williamson) one already possesses. With greater knowledge, new cognitive projects will be available; the cognitive project of investigating the possibility of life on Mars did not have any promise for an ancient Greek, but it has promise for us. This would make sense of Williamson's idea, discussed above, that the prior probabilities will change with contexts; different states of evidence will make different cognitive projects valuable, and will therefore determine different distributions of epistemic probability. Finally note that, although it will be typically the case that a cognitive

[19] It should be stressed that naturalness is not to be identified with simplicity, as the name could perhaps suggest. If it were the proposal would reduce to a fairly familiar one. More importantly, I do not think anyone has been able to make plausible the idea that simplicity, by itself, could play the desired role. The difference in simplicity, on any reasonable understanding of this notion, between a sceptical hypothesis and its negation seems scarcely significant, and surely not big enough to set them apart in the required way.

[20] I am not of course defending externalism here. I simply note that there is a view in the vicinity of Wright's which is compatible with externalist assumptions. I am in fact convinced that a form of pluralism is needed in epistemology with respect to the internalism/externalism divide; therefore, I also think there is intrinsic interest in developing this externalist notion of entitlement. But I am not arguing for this view here.

project with a false presupposition will not produce knowledge, this must not be an absolute rule. Occasionally, a project with a false presupposition might produce accurate results reliably enough for them to constitute knowledge. So, on this version of the view, while the truth of the presupposition is not irrelevant, it is not a necessary condition for the proposition to have a certain degree of entitlement, so to speak. On this picture, the subject is not irrelevant to the epistemic status of the proposition. But (unlike Wright's view) the view I am describing still leaves a large role to the relation of the proposition with external facts. Some propositions, such as the negation of a sceptical hypothesis, will by their content be more suited than others to be presuppositions of a fruitful cognitive project, for subjects like we are anyway, and will therefore tend to attract a high degree of justification.

As I anticipated, these remarks only provided a rough sketch of a possible direction of investigation. It will be for another work (and perhaps, I hope, another author) to put some flesh on these bones.

5 Conclusion

I described two ways in which one might commit the consequence fallacy, the fallacy, roughly, of rejecting a defeasible inference because there is consequence of its conclusion which is not supported by its premises. The first way, which is rather trivial, involves overlooking background information; the second way, which is the more philosophically interesting, involves overlooking prior probabilities. I then discussed the relation between the consequence fallacy and some sceptical argument, in particular in an argument against the possibility of defeasible justification I develop elaborating on previous work by Huemer. I argued that one can resist that form of sceptical argument if one gives up the idea that a belief cannot be justified unless it is supported by the totality of the evidence available to the subject—a principle which is entailed by many prominent epistemological views, most clearly by epistemological evidentialism. The justification of the belief, in the cases for which the principle fails, should derive not from the evidence but rather from the prior probability of the proposition. In other words, I argued that we can resist this kind of sceptical argument provided that prior probabilities yield a form of justification that does not rely on evidence, so an *entitlement*, in Crispin Wright's sense, for the negation of sceptical hypotheses. I concluded with some discussion of how to understand prior probabilities, and how to develop the notion of entitlement, in an externalist, and even "knowledge-first", epistemological framework. In particular, I sketched, without endorsing it, a view on which we have entitlement for propositions which are indispensable presuppositions of a valuable cognitive project, where the value of a cognitive project is a function of the amount of knowledge it would produce.

Acknowledgments Several people helped me to get clearer about the topic discussed in this paper through the years. I would like to thank in particular Jessica Brown, Diego Marconi, Martin Smith, Brian Weatherson, Tim Williamson and Crispin Wright. Wholehearted thanks also go to an anonymous reviewer for *Philosophical Studies*.

References

Burge, T. (2003). Perceptual entitlement. *Philosophy and Phenomenological Research, 67*, 503–548.

Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives, 2*, 91–123.

Comesana, J. (2010). Evidentialist reliabilism. *Nous, 44*, 571–600.

Dodd, D. (2007). Why Williamson should be a sceptic. *Philosophical Quarterly, 57*, 635–649.

Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies, 48*, 15–34.

Field, H. (2000). Apriority as an evaluative notion. In P. Boghossian & C. Peacocke (Eds.), *New essays on the a priori* (pp. 117–149). Oxford: Clarendon Press.

Firth, R. (1978). Are epistemic concepts reducible to ethical concepts? In A. I. Goldman & J. Kim (Eds.), *Values and morals* (pp. 215–229). Dordrecht: Reidel.

Goldman, A. (2009). Williamson on knowledge and evidence. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge* (pp. 73–91). Oxford: Oxford University Press.

Harman, G. (1986). *Change in view*. Cambridge, MA: MIT Press.

Harman, G. (2003). Skepticism and foundations. In S. Luper (ed.) *The skeptics: Contemporary essays* (pp. 1–11). Aldershot: Ashgate.

Huemer, M. (2001). The problem of defeasible justification. *Erkenntnis, 54*, 375–397.

Ichikawa Jenkins, C. (forthcoming). Justification magnetism, forthcoming in *Philosophical studies*. Retrieved on-line on 16/06/2012 at http://www.carriejenkins.co.uk/research/papers.

Lewis, D. K. (1984). Putnam's paradox. *Australasian Journal of Philosophy, 62*, 221–236.

Schiffer, S. (2004). Skepticism and the vagaries of justified belief, *Philosophical Studies, 119*, 161–184.

Shackel, N. (2007). Bertrand's paradox and the principle of indifference. *Philosophy of Science, 74*, 150–175.

Stine, G. C. (1976). Skepticism, relevant alternatives and deductive closure. *Philosophical Studies, 29*, 249–261.

Swinburne, R. (2001). *Epistemic justification*. Oxford: Oxford University Press.

Weatherson, B. (2007). The bayesian and the dogmatist. *Proceedings of the Aristotelian Society, 107*, 169–185.

Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.

Williamson, T. (2007). *The philosophy of philosophy*. Oxford: Blackwell Publishing.

Wong, W., & Zudell, Z. (2012). How fallacious is the consequence fallacy? Forthcoming in *Philosophical studies*, published online 27 April 2012.

Wright, C. (2004). Warrant for nothing (and foundations for free)? *The Aristotelian Society Supplementary Volume, 78*, 167–212.