

Abstract

Williamson (2007) describes what he calls ‘the consequence fallacy’, which consists, roughly, in inferring from the fact that a certain hypothesis has a consequence which is not supported by our evidence, to the conclusion that we are not entitled to a high degree of confidence in either the hypothesis itself or its consequence.

I agree this is a fallacy. The aim of the paper is to clarify its nature, and to assess its significance.

I will propose two different, but not incompatible, explanations of the fallacy. I will then discuss a form of sceptical argument which does commit the fallacy, in one of the two ways. The fallacy consists, in the relevant cases, in overlooking the role of prior probabilities.

In the last section, I will claim however that more work is needed on the notion of prior probability, and that a successful explanation of the consequence fallacy in connection with sceptical arguments requires in particular a notion of entitlement based on prior probability.

Consequence Fallacy, Defeasible Justification and Entitlement

Section 1 The fallacy

In this section, we will describe the “consequence fallacy”, and we will try to understand it properly. The notion is introduced in Williamson (2007) 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 $\text{Prob}(h) < \text{Prob}(h | e) < 1$. The material conditional $e \rightarrow h$ is a logical consequence of h , and therefore at least as probable as h ; in fact, $\text{Prob}(e \rightarrow h | e) = \text{Prob}(h | e)$. However, e is evidence against $e \rightarrow h$, for $\text{Prob}(e \rightarrow h) > \text{Prob}(e \rightarrow h | e)$, simply because $e \rightarrow h$ is true in all those possibilities which e eliminates ($e \rightarrow h$ is a logical consequence of $\neg e$). Clearly, all of this is compatible with a high degree of legitimate confidence in both h and $e \rightarrow h$. 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*.”¹

Here is a rather trivial way of committing the fallacy: 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, if the gauge points to empty then the tank is empty, on the basis of the fact that the gauge

¹ Williamson (2007) p. 232

points to empty. But surely I cannot do that; whatever the gauge indicates is not by itself evidence on its reliability. So I refrain from forming the belief that I am out of gas.

Why would that be poor reasoning? Well, for one thing, 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 that $e \rightarrow h$.

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 a logical consequence of **h**, namely $e \rightarrow h$, 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 $e \rightarrow h$.

So, one would like to ask, how are we entitled to a high degree of confidence in $e \rightarrow h$?

Williamson does have 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**].”²

As I said, Williamson is not explicit about embracing this explanation of the fallacy. 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 requires³ of the sort of probability he is talking about⁴. As Williamson notes, $\text{Prob}(e \rightarrow h | e) = \text{Prob}(h | e)$, and $\text{Prob}(e \rightarrow h) > \text{Prob}(e \rightarrow h | e)$. But, if we are entitled to a high

² Williamson (2007) p. 228

³ Williamson (2007), p. 228, refers back to Williamson (2000), and the latter, p. 211, makes explicit the assumption of standard axioms for probability calculus. Pryor (manuscript) proposes a revision of the standard axioms in the light of a similar problem. I won't discuss Pryor's proposal here.

⁴ We will come back to what this sort of probability is. Note however that the basic axioms seem independent of the interpretation of probability we will opt for

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.

I will come back to this explanation of the consequence fallacy, and in particular to the notion of prior probability.

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 one can reasonably think that there are relations between background beliefs and prior probabilities. Still, it’s important to distinguish the two sorts of cases, for reasons that will become clear while we evaluate the significance of the consequence fallacy.

Section 2. Consequence fallacy and scepticism about defeasible justification

I think the fallacy has a role in many sceptical arguments. I will consider just one⁵, which seems particularly relevant. It is an argument presented, although not endorsed, in Huemer (2001), for the impossibility 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. Since Williamson is assuming all evidence to be propositional, we can put aside this limitation. I will reformulate Huemer’s argument accordingly.

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** and $\neg\mathbf{p}$) is consistent. In such cases Huemer notes there will always be an hypothesis **h** which entails **e** and is incompatible with **p**. Sceptical hypotheses of great imaginative force abound; however, the most simple case is to consider the hypothesis **h** which is simply (**e** and $\neg\mathbf{p}$) . The support given by our evidence should be at least strong enough for the subject to be justified in believing the conclusion; and strong enough, in absence of other obstacles, for that belief to qualify as knowledge.

Huemer’s argument to the effect that this is never going to be the case uses two principles, which I slightly re-formulate, with respect to Huemer’s version, as follows:

⁵ Schiffer (forthcoming) also seems to me to present a form of sceptical argument in which the consequence fallacy is exemplified.

P1 Given any evidence, **e**, and hypothesis **h**, in order for **e** to justify **h**, one must have at least some reason independent of **h** for rejecting each of the alternatives incompatible with **h** which entail **e**.⁶

P2 For any set of propositions **e** and **h**, if **h** entails **e**, then **e** is not a reason for rejecting **h**.⁷

Here is an informal statement of the argument. Take a particular case, e.g, 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** implies **e**, by P2 it is not the case that **e** gives reason for rejecting **sh**; since **sh** is incompatible with **h**, by P1 **e** cannot justify **h**, unless we can rule out **sh** in some other way. Do we have other reasons, independent of **h**, for rejecting **sh**? If we do, let's add those reasons to our initial evidence **e**, and get the new evidence **e***. If **e*** rules out every possibility incompatible with **h**, then we have a 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* which is incompatible with h and entails e | From 1 |
| 3) You do not have reasons independent of h to reject h* | From 1), 2) and P2 |
| 4) e does not justify h | From 2), 3) and P1 |

That the argument exemplifies the consequence fallacy should be quite obvious. It moves from the logical possibility of $e \wedge \neg h$, which is of course equivalent to the logical possibility of the negation of the material conditional $e \rightarrow h$, to the conclusion that **e** does not support **h**. While Williamson noted that endorsing the argument form he dubs "consequence fallacy" would have led to the (absurd) conclusion that no defeasible justification is possible, the argument Huemer presents seems to provide independent plausible reasons to think that the argument form is indeed flawed, and it concludes that no defeasible justification is possible.

⁶ Huemer (2001) had, at the the end of P1 "in which e would be true". I think what was meant was just 'which entails e', and I use this clearer, to my mind, formulation.

⁷ P2 is clearly false, assuming classical logic, where **e** is a logical truth or **h** is a contradiction. I will leave the limitation implicit.

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 no reason to reject it. But, by 1), e is our total evidence; still, one might think, there might be reasons to reject h^* that are 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 they are 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 If your total evidence e is not a reason to reject an hypothesis h , then you have not a reason to reject h .

If one accepts P1 to P3, one has to deny the consequence fallacy is a fallacy, and to embrace the sceptical conclusion as to the possibility of defeasible justification. I think P3 should be rejected. Let us look at alternative ways of locating the error. Since e is stipulated to be our total evidence, background information is not playing a role here. Therefore, it must be a case of ignoring prior probabilities. We face however a choice: leaving aside P1, which seems tantamount to a principle of closure for justification, we have to individuate the principle to be rejected among P2 and P3. However, it seems the denial of P2 could be defended only on a peculiar reading of 'being a reason for'. If h entails e , then e makes h more likely than it was; it might still be the case that one has reasons to reject h , but that is *despite e*. It's true that there are possible evidential states such that the low prior probability of h would have been raised so much that you would not have been entitled to reject any more, but that does not seem enough to say that e is a reason to reject h . Compare the following practical case: Suppose I am told to open the door, and I then proceed to open the window. I then claim to have acted on the basis of what I was told. When someone observes there seems to be no relation between what I was told to do and what I did, I reply that I already had a strong inclination to open the window; therefore, I would have opened it on the basis of a large class of things I was told to do. Not anything would have done; if I had been told to help myself to a piece of cake, e.g., I would not have opened the window. Therefore, I still opened the window on the basis of what I was told. I don't think anyone would, or should, be convinced by my reply. But I don't see any relevant difference in the epistemic case⁸.

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.

⁸ This is relevant to the evaluation of what Williamson calls "appearance principles", which state roughly that one can be confident of a proposition on the basis of a source only if the appearance that the proposition is true, related to that source, is good evidence for it. Williamson thinks the fallacy is connected to acceptance of this kind of principle. I think the sort of reasoning we applied to P2 suggests instead that the principles are compatible with the recognition of the fallacy, and indeed are still plausible even when the fallacy is avoided. I am leaving the point aside for reasons of space.

But we have a better alternative, which is denying P3. P3 implies that prior probabilities cannot provide reasons to reject or accept an hypothesis; but we have seen that whenever a (total) body of evidence supports an hypothesis without entailing it, there is a consequence of the hypothesis which we are entitled to a high degree of confidence about, by its prior probability already.

It is now time to give a better look to the notion of prior probability.

Section 3. 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.”⁹ This is almost everything which is said about prior probabilities, except that “the characterization of the prior distribution of evidential probabilities is blatantly vague”.¹⁰

This of course is not incredibly encouraging. One interesting point which I will have to leave apart is Williamson’s remark on the variation between “contexts”; I am not clear 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, at this point I am not clear enough on what the proposal would look like, and I have to leave it aside.

Something more needs to be said on the notion of prior probability. 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, where the totality of our evidence is considered, and which respects standard axioms of 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 we are anyway entitled to a high degree of confidence about. This seems to me to have the 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

⁹ Williamson (2000) p. 211

¹⁰ Williamson (2000) p. 212

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.

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.”¹¹

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. 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. Both points are well illustrated by the present problem. I cannot argue for this negative point here. Let me just note about the most prominent alternative, the subjective interpretation, that it 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.

Even if the negative point is granted, however, it only makes the situation worse, as no well worked-out alternative notion is in sight. Williamson observes that the lack of “an operational definition” is not a good reason to avoid using the notion. I certainly agree much of worth 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 (2004) notion of entitlement; a form of warrant which comes “for free”, independently of the evidence.

Let me clarify the latter claim going through an abstract reconstruction of a kind of reasoning which may lie behind the consequence fallacy:

- 1) Suppose on the basis of **e** I acquired justification to believe **h**, although **e** does not entail **h**
- 2) Since $\mathbf{e} \rightarrow \mathbf{h}$ is a consequence of **h**, I must be now able to infer $\mathbf{e} \rightarrow \mathbf{h}$ from **h**, thus acquiring justification for $\mathbf{e} \rightarrow \mathbf{h}$.

¹¹ Williamson (2007) p 228

3) My belief in **h** is supported by **e**, so if I acquired a justified belief in $e \rightarrow h$ through inferring it from **h**, my belief in $e \rightarrow h$ would be supported by **e**

4) Since **e** lowers the probability of $e \rightarrow h$, it cannot support $e \rightarrow h$

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; 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 $e \rightarrow h$, then I have justification to believe $e \rightarrow h$; it does not demand that I am able to acquire a justification for $e \rightarrow h$ through **h**, because to be justified in believing **h** might require to be already justified in believing $e \rightarrow h$ in the first place. But, supposing that it is not impossible to acquire justification (as opposed to mere high degree of confidence) through defeasible evidence, what the prior probability has to provide for the consequence is nothing less than justification.

Conclusion

.We described two ways in which one might commit the consequence fallacy. The first, which is rather trivial, involves overlooking background information; the second, more interesting, way involves overlooking prior probabilities. We have then seen the consequence fallacy at work in an argument against the possibility of defeasible justification. Recognizing it can help with this form of argument, provided we have a satisfying understanding of prior probabilities, and of a form of warrant based solely on those.

References

- Huemer, M., 2001, "The Problem of Defeasible Justification", *Erkenntnis*, 54, 375-97
- Pryor, J. , manuscript, "Uncertainty and Undermining"
- Schiffer, S, forthcoming "Evidence = Knowledge: Williamson's Solution to Skepticism " , forthcoming in P. Greenough and D. Pritchard, eds., *Williamson on Knowledge*
- Williamson, T., 2000, *Knowledge and its Limits*, Oxford, Oxford University Press
- 2007, *The Philosophy of Philosophy*, Oxford, Oxford University Press
- Wright, C., 2004, "Warrant for nothing (and foundations for free)?", *The Aristotelian Society Supplementary Volume*, 78, 167-212

