

Williamson on Knowledge and Evidence

Alvin Goldman
Rutgers University

[To appear in D. Pritchard, ed., Williamson on Knowledge, OUP]

1. Methodological Questions

Timothy Williamson's project in Knowledge and Its Limits (Williamson, 2000)¹ includes proposals for substantial revisions in the received approach to epistemology. One received view is that knowledge is conceptualized in terms of a conjunction of factors that are individually necessary and jointly sufficient for knowing. A central aim of epistemology is to state such necessary and sufficient conditions. Against this received view, Williamson argues that

a necessary but insufficient condition need not be a conjunct of a non-circular necessary and sufficient condition. Although being coloured is a necessary but insufficient condition for being red, we cannot state a necessary and sufficient condition for being red by conjoining being coloured with other properties specified without reference to red. Neither the equation 'Red = coloured + X' nor the equation 'Knowledge = true belief + X' need have a non-circular solution. (3)

Williamson further argues that we have inductive reasons for thinking that no analysis of the concept knows of the standard kind is correct. The inductive reasons are simply the history of failed attempts at such "factorizing" or "decompositional" analyses. Williamson not only rejects the prospect of explaining knowledge in terms of belief, justification, and evidence, but he proposes to reverse the order of explanation.

That order of explanation has been reversed in this book. The concept knows is fundamental, the primary implement of epistemological inquiry. (185)

It is not altogether easy, however, to reconcile this radical program with other things Williamson says in the book. In particular, the book contains two rather unrelated accounts of knowing, and one of these accounts appeals to some of the same ingredients that traditional (or semi-traditional) analysts have used in the past. The first account, which appears in Chapter 1, says that knowing is the most general factive stative propositional attitude. The account is presented in terms of three conditions:

- (1) If Φ is an FMSO [factive mental state operator], from 'S Φ s that A' one may infer 'A'.
- (2) 'Know' is an FMSO.
- (3) If Φ is an FMSO, from 'S Φ s that A' one may infer 'S knows that A'. (2000: 39)

The latter two principles, says Williamson, “characterize the concept of knowing uniquely, up to logical equivalence, in terms of the concept of an FMSO” (39). Williamson adds that this account is not a decomposition of the knowledge concept; it would be implausible to claim that everyone who thinks that John knows that it is raining thereby thinks that John has the most general factive stative propositional attitude to the proposition that it is raining. This remark is intended, no doubt, to distinguish this characterization of knowledge from traditional analyses.

But the second type of account Williamson offers seems closer to traditional analyses. He acknowledges that knowing seems to be “highly sensitive” to such factors as justification, causation, and reliability, over wide ranges of cases. “Any adequate account of knowing should enable one to understand these connections” (41, emphasis added). He clearly indicates that he regards truth, belief, and reliability (more precisely, safety) as necessary conditions for knowing. “No reason has emerged to doubt the intuitive claim that reliability is necessary for knowledge” (100), and “A reliability condition on knowledge was implicit in the argument of section 5.1 and explicit in sections 4.3 and 4.4” (123). So Williamson does not abandon the prospect of giving an account, if not an analysis, of knowing, along something approaching traditional lines. Such an enterprise is sprinkled through the middle chapters of the book. Although his account never formulates sufficient conditions for knowing, and therein departs from the enterprise of traditional analysis, it does go some distance in a fairly orthodox direction. It is therefore a bit difficult to piece together the different strands of his methodological thinking.

I wish to examine the “second” account (or partial account) of knowledge Williamson offers. In light of his rejection of traditional analyses, however, it is unclear how to proceed. What criteria of adequacy should be imposed on an account of the sort Williamson offers? What standards should be used in judging this attempt? Williamson himself, of course, denies that an account is required to provide a sufficient and non-circular condition for knowing, and he doesn’t attempt to provide one. But what if we are not yet persuaded that knowing has no non-circular sufficient condition (of the traditional kind)? How should we rate the success of Williamson’s account compared to rival accounts that do offer sufficient (and non-circular) conditions?

To illustrate the puzzle, return to an epistemological “yesterday,” when Gettier (1963) first propounded his counterexamples to the justified true belief (JTB) analysis of knowing. Suppose that supporters of the JTB analysis had responded: “Professor Gettier, your two examples are extremely interesting. But what they purport to show is that JTB is not sufficient for knowing. We, however, do not claim that JTB is a sufficient condition; we don’t believe that the knowledge concept admits of any non-circular sufficient condition. So the JTB account is completely in order as it stands. It is wholly adequate, and not open to the kind of critique via non-sufficiency that your examples are intended to support.”

Few epistemologists would be impressed by this response, either then or now; nor should they be impressed. Whether or not Williamson is right about the

“unanalyzability” of knows, the JTB analysis obviously misses some important facts or elements, as post-Gettier investigations have revealed. To rest content with the JTB account would be a lauder today. It is obviously incomplete, and we have some fairly good ideas about how to repair it, or at least improve upon it. This is not to say there is unanimity among epistemologists. Nevertheless, some sort of additional conditions in an “anti-luck” vein are widely agreed to be necessary for a satisfactory account of knowing. Thus, the procedure of testing a set of conditions for sufficiency, beautifully exemplified by Gettier, is a salutary epistemological exercise.

Let me turn this example into a general proposal for an adequacy condition on accounts of knowing (or other philosophically interesting concepts). A fully adequate set of conditions, I propose, should be both correct and complete. Completeness, however, can be interpreted in different ways, depending on whether the target of the account is or isn't susceptible of (non-circular) sufficient conditions. If sufficient conditions exist, then a complete account should provide all necessary and sufficient conditions (that are philosophically significant). If no (non-circular) sufficient conditions exist, as in the case of red, then a complete account should provide all necessary conditions (that are philosophically significant). Under this proposal, even someone like Williamson who is dubious about “factorizing” analyses could agree that Gettier proved something important, namely, that the JTB account omits some (non-circular) necessary conditions for knowing, and hence is not fully adequate.²

The moral I would extract from this discussion is the following. It always pays to subject any proposed account of a concept to sufficiency tests. If an account fails a sufficiency test, this is compatible with any of three alternative situations. First, there may be some non-circular sufficient condition that hasn't yet been identified. Second, there may be one or more necessary (but insufficient) conditions that have been omitted and remain to be identified. Third, there may be no further non-circular conditions, even necessary conditions, which could be added. Perhaps all additional necessary conditions are circular ones. Failing a sufficiency test doesn't by itself establish which of these three alternatives obtains. Nonetheless, discovery of a sufficiency failure signals the need for further investigation into the concept, because progress is possible if either the first or second alternative obtains. Hence, subjecting proposed accounts to sufficiency tests is a desirable diagnostic technique.

This is the moral I shall pursue in some of the examination of Williamson's account of knowing that follows. I shall sometimes ask whether his account of knowing passes a sufficiency test. I pose this question despite the fact that Williamson himself doesn't claim sufficiency for his conditions. Still, this is a question that should interest the epistemological community, for the reasons given above. Furthermore, since Williamson proposes the principle that evidence is co-extensive with knowledge (see section 6 below), it would be good to pinpoint all conditions required for a state to qualify as knowledge, not just some of them. Before agreeing to the E = K thesis, it would be good to know everything implied by knowing. If Williamson has given us an incomplete list of conditions for knowing, we should be aware of this, lest we adopt a conception of evidence without fully appreciating what it commits us to.

Let me clarify what my little critique of Williamson's account of knowledge won't be aimed at. Although I shall question the adequacy of his account, I won't challenge any of the specific implications he draws from this account. A major part of Williamson's aim is to derive certain theoretical conclusions from a safety account of knowing, such conclusions as anti-luminosity and the denial of the KK principle. These conclusions are drawn from the necessity of the conditions he proposes. I won't take issue with any of these conclusions.³ I suspect that any changes in the account of knowing that my worries might motivate would leave these implications intact. Nonetheless, epistemologists are interested in getting as adequate an account of knowing as possible. My objections are devoted to this end.

2. The Non-Sufficiency of Safety

The account of knowing Williamson develops in chapters 4-7 belongs to the reliability family. Williamson locates his preferred form of reliability in the "safety" category, but there remain questions about the account's adequacy and its alleged superiority to other theories in the reliability family.

Williamson's first pass at the kind of reliability he regards as relevant to knowledge occurs in the following passage:

For present purposes we are interested in a notion of reliability on which, in given circumstances, something happens reliably if and only if it is not in danger of not happening. That is, it happens reliably in a case α if and only if it happens (reliably or not) in every case similar enough to α . In particular, one avoids false belief reliably in α if and only if one avoids false belief in every case similar enough to α . (124)

Elsewhere he refers to this reliability notion as "safety from error," and presents a principle for knowledge that invokes the safety-from-error notion.

Now assume a connection between knowledge and safety from error: ... For all cases α and β , if β is close to α and in α one knows that C obtains, then in β one does not falsely believe that C obtains. (128)

Here we have, in kernel form, Williamson's safety account of knowing. There are two straightforward questions to ask about it. Is safety necessary for knowing, and is it sufficient? (Of course, Williamson doesn't claim that safety is sufficient. But since he offers no further condition, we have to ask whether the condition he has offered is "complete.") The sufficiency question is addressed in this section and section 5, and the necessity question is addressed in sections 3 and 4.

Sherrilyn Roush (2005: 122-123) presents a counterexample to the sufficiency of the safety account -- not Williamson's version of safety, exactly, but near enough. S has a fairy godmother whose special mode of operation is to make true anything that S

believes. So, for any p , when S actually believes p , it is true. Moreover, in close counterfactual situations, the fairy godmother would still make p true if S believed it. So the safety condition is fulfilled. But fulfillment of this condition doesn't guarantee that S 's belief in p intuitively counts as knowing. To flesh out an appropriate example, suppose S comes to believe p by applying a fallacious mode of reasoning to some false and unjustified prior beliefs q and r . Or, suppose S comes to believe p by idle wishful thinking. S lacks even a wisp of evidence, or what she takes to be evidence, for p . Still she believes it, and the fairy godmother secures its truth. Does S 's belief in p qualify as knowledge? Intuitively, no, though the safety condition is satisfied.

Is the safety requirement really satisfied, one might wonder? Isn't there a close situation in which the fairy godmother takes a break from her duties, or simply doesn't exist? In such a situation, p might easily be false though S still believes it.⁴ However, we can stipulate facts about the actual case that make the absence or dereliction of the fairy godmother in a counterfactual case extremely remote, hence not a threat to safety. Suppose the actual world is one in which certain human agents, including S , have a fairy godmother as a matter of nomological necessity. Then the non-existence of such a godmother doesn't obtain in any close world. The nomological necessity version of the case allows metaphysically possible worlds in which there is no fairy godmother; such worlds are just extremely remote.

A different objection to this example is that if S has a fairy godmother who always ensures the truth of S 's beliefs, at least for future-oriented beliefs, then S will accumulate excellent inductive evidence that her future-oriented beliefs always turn out true. If so, won't S know each future-oriented proposition she believes?⁵ Here we need to be careful about the time of the belief. Until S has time to reflect on the fact of her believing, say, p , she doesn't yet have inductive evidence for that particular proposition. Only after the belief in p is formed can S identify it as one of her beliefs and infer from her believing it that it will somehow come true. Until this further inference is made, her belief doesn't qualify as knowledge. But the belief is safe even before the inferential step is taken, which shows that safety doesn't suffice for knowing.

There are other kinds of cases that can also serve as counterexamples to the sufficiency of safety. Suppose S is looking at an orange harvest moon perched over the horizon (Goldman 1976). S correctly believes that it is the moon, and in possible cases in which the moon has slightly different properties, he would still believe it is the moon. For example, if it were slightly larger or smaller, if its rocky face had slightly different peaks or valleys, etc., S would still correctly believe it is the moon. In short, S avoids false belief in every case similar enough to the actual one; or so it initially seems.

Now consider a counterfactual scenario, one in which what is perched over the horizon isn't the moon but a large orange hot-air balloon, which looks from S 's perspective just as the moon looks to him in the actual scenario. If this case transpired, S would mistake the hot-air balloon for the moon, i.e., would mistakenly believe that the perceived object is the moon. Of course, the large orange balloon is vastly closer to S than the moon is, and in physical terms – size, composition, etc. – bears little

resemblance to the moon. But it would present a very similar appearance to S as the moon presents in the actual scenario. Finally, assume that orange hot-air balloons are moderately common in S's part of the world. Then I am inclined to say that S doesn't know, in the actual situation, that what he sees is the moon.

Does Williamson's safety condition handle this case correctly? Does S's moon belief pass Williamson's safety test and thereby qualify for knowledge? If so, the safety condition isn't sufficient for knowing. So let's ask whether the moon belief passes the safety test.

Williamson says that his notions of reliability, stability, safety, and robustness concern what could easily have happened. They depend on what happens under "small variations in initial conditions" (123). When illustrating these ideas, Williamson invariably chooses examples featuring small variations in the physical features of the objects in question, e.g., small variations in the height of a tree. If we take this to be a prototype of his notion of 'closeness,' it appears that the case in which the moon is present in S's visual field and the case in which a hot-air balloon is present in S's visual field are extremely distant from one another. The moon is vastly larger than a hot-air balloon, has vastly greater mass, is composed of very different material, and is much farther from S than is the hot-air balloon. So the balloon case wouldn't seem to qualify as 'close' to the moon case. Hence, it shouldn't matter under the safety requirement that S makes a mistake in the balloon case. S's belief in the moon case appears to be safe, and therefore should be an item of knowledge according to the safety account. But; intuitively, the balloon alternative keeps S from knowing. Thus, passing the safety condition isn't sufficient for knowing.

Perhaps Williamson's criterion of closeness should be interpreted differently. Maybe his notion of closeness is intended to incorporate relativization to a cognizer's evidence in the actual state of affairs. Applying this idea to the present case, the moon/balloon scenarios would be close to one another relative to S's perceptual experience. Williamson, however, never incorporates relativization-to-experience into his account of closeness; this is my suggestion. Moreover, it is doubtful that Williamson would accept the proposed role for 'experience'. He does propose to place a 'sameness of basis' condition on his account of safety (see section 3 below), but sameness of 'experience' threatens to introduce an internal conception of a 'basis', which Williamson rejects. In any case, even a guarantee of experiential equivalence doesn't prove that the balloon scenario is a close enough alternative to the moon scenario to disqualify the latter as an instance of knowledge.

3. Safety as Necessary for Knowing

Is safety, as characterized by Williamson, necessary for knowing? Neta and Rohrbaugh (2004) offer two counterexamples to its necessity. Here is their presentation of the examples:

- (A) I am drinking a glass of water which I have just poured from the bottle. Standing next to me is a happy person who has just won the lottery. Had this person lost the lottery, she would have maliciously polluted my water with a tasteless, odorless, colorless toxin. But since she won the lottery, she does no such thing. Nonetheless, she almost lost the lottery. Now, I drink the pure, unadulterated water and judge, truly and knowingly, that I am drinking pure, unadulterated water. But the toxin would not have flavored the water, and so had the toxin gone in, I would still have believed falsely that I was drinking pure, unadulterated water.... Despite the falsity of my belief in the nearby possibility, it seems that, in the actual case, I know that I am drinking pure, unadulterated water.
- (B) I am participating in a psychological experiment, in which I am to report the number of flashes I recall being shown. Before being shown the stimuli, I consume a glass of liquid at the request of the experimenter. Unbeknownst to either of us, I have been randomly assigned to the control group, and the glass contains ordinary orange juice. Other experimental groups receive juice mixed with one of a variety of chemicals which hinder the functioning of memory without a detectable phenomenological difference. I am shown seven flashes and judge, truly and knowingly, that I have been shown seven flashes. Had I been a member of one of the experimental groups to which I was almost assigned, I would have been shown only six flashes but still believed that I had been shown seven flashes due to the effects of the drug. It seems that in the actual case I know that the number of flashes is seven despite the envisaged possibility of my being wrong. And yet these possibilities are as similar in other respects as they would have to be for the experiment to be well designed and properly executed. (2004: 399-400)

Neta and Rohrbaugh's examples strike me, intuitively, as correct. However, they rightly consider the possible objection that their examples shouldn't be considered cases of knowledge because of their similarity to the fake barns example (Goldman 1976, although the example is originally due to Carl Ginet⁶). Most people agree that in the fake barns example Henry doesn't know he is seeing a barn (though he is), apparently because of the close possibility that what he is seeing is a mere façade. Neta and Rohrbaugh argue that their examples (A) and (B) aren't analogous to the fake barn case because Henry's actual circumstances are unfavorable in a way in which the actual circumstances of the agents of (A) and (B) are not. There really are fake barns in Henry's neighborhood. By contrast, the threats to knowledge in (A) and (B) remain purely counterfactual. Although things could have gone epistemically less well, and almost did so, in each case the threat was avoided and the actual case remains epistemically unproblematic. This response is plausible but not wholly compelling. After all, one might say in the fake barns case that the threat of a fake barn being before Henry remains counterfactual, so the threat to Henry's knowledge, if Neta and Rohrbaugh are right, ought to be avoided.

However, the problem should not be laid at Neta and Rohrbaugh's doorstep. There remains the important question that should be directed to Williamson, viz., how his

theory intends to interpret “closeness”? The question can be divided into two. First, which “qualitative” features of objects in a pair of cases α and β are relevant to determining their degree of closeness? Second, which kinds of world transformations count as staying relatively close to actuality versus straying farther afield? I would expect Williamson to respond by saying that all these matters are simply vague (see p. 100). I don’t necessarily disagree with Williamson about the appropriateness of this response. But it does raise a question of whether his theory constitutes any advance over an earlier theory in the reliabilist tradition, viz., the relevant-alternatives theory. If we interpret a ‘relevant’ alternative as a ‘close’ alternative, then the relevant-alternatives theory also says that a true belief fails to be knowledge if there is a close alternative in which the agent does not avoid having a false belief. I shall return to a comparison of Williamson’s theory and relevant alternatives theory in the next section.

Let us further probe the necessity of the safety condition and the interpretation of closeness by looking at the dachshund-wolf example (Goldman 1976). Depending on how closeness is interpreted, this threatens to be a counterexample to the safety condition. I look at a nearby dachshund and truly believe that what I see is a dog. Had I not been seeing a dachshund, I would have been seeing a wolf, and would have falsely believed myself to be seeing a dog. Clearly, this is a case of knowing. At any rate, the fact that what I would see if the dachshund weren’t there is a wolf, and I would mistake the wolf for a dog, doesn’t disqualify this as an instance of knowing. Intuitively, this is because the wolf wouldn’t look anything like the dachshund. However, my belief in the actual case that what I am seeing is a dog may not be safe, because the wolf case may be close enough to the dachshund case. Why would it be close? First, the wolf case is close to the dachshund case because the world-transformation, as we stipulated, is a minor one. Second, the wolf physically resembles the dachshund to a moderate degree. They are both furry quadrupedal mammals belonging to the canine family. If Williamson chose to weight the dimension of world-transformation over the dimension of physical similarity, this case could well be in one in which there is knowledge but the safety condition is violated.

Williamson discusses the dachshund-wolf example in his critique of Nozick’s (1981) sensitivity theory, but he is silent about how his own theory would handle the example. He might proceed by appealing to the fact that the basis for belief in the wolf case is different from the basis for belief in the dachshund case. While discussing Nozick, who invokes sameness of ‘method’ in his sensitivity theory, Williamson indicates that he favors an improvement in his own safety theory that introduces a sameness-of-basis qualification. Recall that his earlier formulation of safety was: “For all cases α and β , if β is close to α and in α one knows that C obtains, then in β one does not falsely believe that C obtains.” He now writes:

In a more careful version of [the foregoing], we might qualify both ‘know’ and ‘believe’ by ‘on a basis B’. Knowledge on one basis (for example, seeing an event) is quite consistent with false belief in a close case on a very different basis (for example, hearing about the event). (128)

With the addition of a sameness-of-basis qualification, Williamson might argue that knowledge in the dachshund case isn't disqualified because belief in the wolf case does not have the same basis as belief in the dachshund case.

Whether this argument succeeds depends on how 'bases' are construed. As we noted in section 2, one possibility is to construe bases, at least in perceptual examples, as perceptual experiences or appearances. Since the wolf appearance is quite different from the dachshund appearance, the sameness-of-belief requirement wouldn't be fulfilled. Hence the safety condition isn't violated, yielding the verdict that the agent knows in the dachshund case, in conformity with intuition. However, it is most unlikely that Williamson would avail himself of this approach, because appearances are likely to be individuated internally, and he staunchly opposes internal individuation of bases. In discussing Nozick's theory, Williamson cites Nozick's endorsement of an internal criterion of method individuation, and expresses opposition to that criterion.

Another possibility is to construe 'bases' as very general methods for arriving at belief. In the passage quoted above, the examples of 'bases' are seeing and hearing. The general-methods interpretation of 'bases,' however, won't help with the dachshund-wolf case, because the agent uses vision in both the dachshund and the wolf case. So the sameness-of-basis test is passed under this very general interpretation of 'bases', which implies that safety is violated in the dachshund case. Yet, intuitively, the agent knows.

A third possible construal of 'bases' would include specific external objects involved in the method of belief acquisition. The basis of belief in the dachshund case might be seeing the dachshund, and the basis of belief in the wolf case might be seeing the wolf. This kind of external individuation would suit Williamson's favored approach quite nicely. And it would have the desired result in the dachshund-wolf example, because the two cases would then have different bases. Safety wouldn't be violated, and knowledge would be preserved.

However, this approach doesn't work in the moon-balloon example. Seeing the moon in the actual case and seeing the balloon in the counterfactual case would be different bases. Hence safety wouldn't be violated, and the agent would know it's the moon. But our intuitive verdict about the moon case is that the person doesn't know. Similarly, in an identical-twin case, one shouldn't be said to know that it's Judy when one would mistake Judy's twin Trudy for Judy. But the belief that it's Judy in both cases would have different bases under the present 'externalist' construal of 'basis,' because in the one case the basis is seeing Judy and in the other it's seeing Trudy. Difference of basis implies that safety isn't violated and knowledge is (wrongly) sustained. So under this kind of external individuation of basis, the safety theory doesn't get things right.⁷

4. Safety versus Relevant Alternatives

We noted in the previous section that Williamson's safety theory bears a close resemblance to the relevant alternatives (RA) theory, and we wondered whether the

safety theory constitutes an advance over RA theory. Let us now pursue that question further.

My own version of RA theory (Goldman 1976) added an element not featured in the basic formulation of safety theory. That RA theory (which was restricted to perceptual knowledge) said that S's true belief of object O that it has property F gets disqualified from being knowledge if there is a relevant perceptually equivalent situation in which the perceptual object lacks F (so that S's belief of the perceptual object that it has F would be false). Perceptual equivalence was intended to be understood in an internalist fashion. More precisely, although perceptual equivalence applies to external objects or situations, it does so in virtue of producing a mental condition that supervenes on the physical state of the agent's brain (i.e., the 'percept' experienced by the cognizer). Presumably, Williamson would not endorse this conception of perceptual equivalence as a condition on knowledge. However, as we have seen, Williamson does add a sameness-of-basis qualification to his safety theory. This addition may have been motivated by a desire to handle such examples as the dachshund-wolf example. But we were unable to identify a satisfactory interpretation of 'basis' that suits Williamson's externalist approach and handles pertinent cases properly. So there is reason to prefer the RA theory to his safety theory. At a minimum, there is no reason to regard safety theory as superior to RA theory.

In discussing versions of reliability theory different from his own, especially Nozick's sensitivity theory, Williamson offers an argument against a theory that uses internal criteria of method- or basis-individuation. Let us examine this argument.

If methods are individuated internally, so that whether one is using method M supervenes on the physical state of one's brain, then (3) will indeed have some skeptical consequences. [Formula (3) reads: "Necessarily, if S knows p via method M then if p were false, S would not believe p via M."] But why should one accept (3) on those terms? The internal individuation of methods appears gerrymandered precisely to make trouble for our claims to knowledge of the external world. Moreover, (3) is implausible in some examples when methods are individuated internally. My knowing by sight in the ordinary way that a mountain stands here seems compatible with the assumption that if no mountain had stood here, a bizarre chain of circumstances would have occurred as a result of which I would have hallucinated a mountain of just this appearance. That type of hallucination occurs only in worlds very unlike the actual world, we may suppose, and the mechanism that produces it is absent from the actual world. I actually satisfy (3) for knowing by sight many other things about my present environment, including that there is an icy mountain here; my eyesight is actually working perfectly and I have every ordinary reason to believe that it is. To block the unwarranted consequence of (3) that I do not know that a mountain stands here, one must individuate methods externally rather than internally. (156)

This argument against internal individuation of methods, however, is less than compelling. The argument adduces trouble for internal individuation only when internal

individuation is paired with a sensitivity condition, which allows a very remote possibility to disqualify knowledge. But the prime culprit in producing this inappropriate verdict seems to be the sensitivity approach, not the internal individuation of methods. Williamson pins the blame on internal individuation, but this choice is inadequately motivated. Why not opt for (my version of) RA theory with its internal individuation of bases, rather than Williamson's safety theory?

5. Process Reliability

Having argued that Williamson's brand of safety offers no advance over RA theory, and may even be inferior to it, I now want to compare Williamson's safety theory to another variant of reliabilism. One component of many reliability theories that apparently has little attraction for Williamson is process reliability. The idea is that the global reliability (truth ratio) of the process or method by which a belief is formed (or sustained) is critical to its knowledge status. Both the safety theory and RA theory highlight conditions that mainly concern the modal neighborhood of the target belief. They don't straightforwardly address the global reliability of the psychological process(es) that is (are) causally responsible for the belief. But any theory that ignores the mode-of-causation issue won't successfully capture all relevant conditions on knowing.⁸

Does Williamson's attention to the 'basis' of belief reflect an interest in process reliability? That is far from clear. Although Williamson requires that the belief basis be the same if a counterfactual case is to be relevant to safety, he doesn't say that the basis of the actual belief must be globally reliable. Nor is it clear that what he means by a 'basis' is the sort of thing that could have a global reliability property. Let us distinguish between two ontological types that could bear causal relations to (resultant) beliefs. The first ontological type is a ground for belief, where prime examples of grounds are mental states such as perceptual experiences or (other) beliefs. The second ontological type is a cognitive or computational process or operation, which might take one or more grounds as inputs and produce a belief as an output. It is the second ontological type that has global reliability properties, because the same process or operation is presumably applied to many sets of inputs and generates many belief outputs. Some of these outputs would be true and others false; so the generating process or operation will have a statistical property of generating truths a certain proportion of the time. Instances of the first ontological type – grounds – are not the sorts of things that have global reliability properties. The term 'basis' is most naturally applied to grounds, not to processes or operations. So it seems unlikely that Williamson's appeal to 'bases' in his safety theory is a move toward process reliabilism.⁹

Williamson aside, there could be a safety theory of knowing that appeals to grounds without invoking processes or their reliability. Could such a safety theory be adequate? I argue to the contrary. Suppose S forms a belief in a true proposition p based on inference from a rich set of prior beliefs, beliefs that constitute knowledge (and hence, on Williamson's view, evidence). Further suppose that there are no beliefs in S's total corpus that (justificationally) undermine the premises from which S's belief in p is

generated. Thus, the total set of grounds, or evidential beliefs, on which S bases his belief in p lends strong support to p, and no other evidence possessed by S defeats this support. Finally, suppose there is no close situation to the actual one in which p is false and S believes it on the same basis (i.e., on the same grounds). In other words, S's belief in p passes the safety test. Would it follow from all this that S's belief in p is knowledge?

No. This isn't implied, because there aren't adequate constraints on how S bases his belief on p. The mere fact that the belief is based on the indicated grounds leaves it wide open how S's reasoning proceeds from the grounds to the belief in p. Not all such reasoning processes or methods would license a knowledge attribution for the conclusion. Roughly, a proper method of reasoning is one that properly exploits or reflects the pertinent (e.g., logical, probabilistic, abductive) relationships between the evidence propositions. Such a method would tend to have comparatively strong reliability properties.¹⁰ A poor or improper method of reasoning would be one that fails to take due account of the pertinent relationships between the evidence propositions. It might proceed, for example, in a very random, haphazard, or arbitrary fashion. Such a method might occasionally coincide in its selected output belief with that of an ideal method. It might occasionally produce a belief that is indeed supported by the evidence. But that doesn't mean that it would be a good method in general. On the contrary, its global reliability score would be low. My claim, then, is that if S uses a method of this inferior sort, the resulting belief (in p) would not merit the status of knowledge.¹¹

Would a globally unreliable method of the sort just described automatically lead to violation of the safety condition? Equivalently, does safety imply global reliability? If so, then adding a global-reliability condition would be redundant. However, I don't think that a (local) safety condition does imply global reliability of the method used. Notice that in counterfactual cases covered by the safety condition – i.e., in “close” counterfactual cases -- the items of evidence S uses are assumed to be true (known). Since they are S's basis for belief in p, and the basis is held fixed in counterfactual cases, these evidence propositions are still true in these possible scenarios. So whatever the prospects for p's falsity in such counterfactual situations, they aren't affected by the method S uses to form a belief in p from that specified basis. If we assume in the present case that safety is satisfied – and that has been our assumption – it leaves open the question of whether the belief-forming method used is superior or inferior, i.e., globally reliable or unreliable. I claim that unless a superior, globally reliable, method is used, knowledge is not attributable. That is something Williamson's account neglects. It's an independent necessary condition that should be added.¹²

6. Is Evidence Co-Extensive with Knowledge?

The rest of the paper addresses Williamson's theory of evidence. Williamson contends that a person's body of evidence is all and only what she knows. In other words, $E = K$. Is this thesis correct? Does it provide the best way to think and talk about evidence? Now there are places in his book where Williamson seems to favor the $E = K$ thesis because it fits with his “big picture”, the so-called “knowledge first” programme. But for those of us not (yet) persuaded of this programme, this appeal isn't terribly

helpful. On the other hand, there are places in the book where he gives entirely independent arguments for the E = K thesis. Specifically, section 9.6 of Knowledge and Its Limits (200-203) offers a number of arguments why propositional evidence is always knowledge. Let us review these arguments in detail. Since many of these arguments have the form, “It is hard to see what rival construal of evidence could account for the intuitive plausibility of datum X as well as E = K does,” let me tentatively propose a competing construal of evidence, to test it against the various intuitive phenomena, or “data,” that Williamson adduces.¹³ The rival construal of evidence I propose is:

(NPJ) Proposition P is an item of evidence for S at time t =_{df.} P is non-inferentially, propositionally justified for S at t.

Call this the non-inferential propositional justification (NPJ) interpretation of evidence. By ‘propositional justification’ I mean the sort of justification an agent has vis-à-vis a proposition when she is justified in believing it, whether or not she actually believes it. A standard illustration of non-inferential propositional justification is having an experience as of seeing a computer screen before you. Being in this visual state might make you (prima facie) justified in believing the proposition “There is a computer screen before me,” whether or not you do believe it. On the present account, a person in this condition would have the indicated proposition as an item of evidence.¹⁴

Now let us examine the arguments for E = K that Williamson offers in section 9.6. The first argument runs as follows:

When we prefer an hypothesis h to an hypothesis h* because h explains our evidence e better than h* does, we are standardly assuming e to be known; if we do not know e, why should h’s capacity to explain e confirm h for us? (200)

Our rival thesis, E = NPJ, would account for this phenomenon as follows. When we are propositionally justified in believing a proposition, we commonly (though not inevitably) believe we are so justified. If we believe we are so justified, we will also tend to assume that the proposition is true. If e is true, and if h explains its truth better than h* explains it, then h’s capacity to explain e confirms h for us. There is no need to postulate E = K to explain this phenomenon.

The next two arguments run as follows:

It is likewise hard to see why the probability of h on e should regulate our degree of belief in h unless we know e. Again, an incompatibility between h and e does not rule out h unless e is known. (200)

First, if we believe e to be true (given that we are justified vis-à-vis e), it is natural to regulate our degree of belief in h as a function of the truth of e. Specifically, it is reasonable to regulate our degree of belief in h by the probability of h on e. Second, e doesn’t have to be known (with all the baggage knowledge entails) for an incompatibility

between \underline{h} and \underline{e} to rule out \underline{h} . It suffices for \underline{e} to be true. If \underline{e} entails the falsity of \underline{h} , and \underline{e} is a given (i.e., it's a given that it's true), then \underline{h} can be ruled out.

Next Williamson adduces an extended example from which he derives several arguments for $E = K$.

Suppose that balls are drawn from a bag, with replacement. ... [A]ssume that someone else has already made the draws; I watch them on film. For a suitable number \underline{n} , the following situation can arise. I have seen draws 1 to \underline{n} ; each was red (produced a red ball). I have not yet seen draw $\underline{n}+1$. I reason probabilistically, and form a justified belief that draw $\underline{n}+1$ was red too. My belief is in fact true. But I do not know that draw $\underline{n}+1$ was red. Consider two false hypotheses:

- \underline{h} : Draws 1 to \underline{n} were red; draw $\underline{n}+1$ was black.
- \underline{h}^* : Draw 1 was black; draws 2 to $\underline{n}+1$ were red.

It is natural to say that \underline{h} is consistent with my evidence and that \underline{h}^* is not. In particular, it is consistent with my evidence that draw $\underline{n}+1$ was black; it is not consistent with my evidence that draw 1 was black. Thus my evidence does not include the proposition that draw $\underline{n}+1$ was red. Why not? After all, by hypothesis I have a justified true belief that it was red. The obvious answer is that I do not know that draw $\underline{n}+1$ was red; the unsatisfied necessary condition for evidence is knowledge. (200-201)

An alternative explanation is readily offered by $E = NPJ$. My evidence does not include the proposition that draw $\underline{n}+1$ was red because, although that proposition is justified, it isn't non-inferentially justified. On the contrary, it is justified by inductive inference. So the unsatisfied necessary condition for evidence, according to $E = NPJ$, is that evidence confers non-inferential justification.

Williamson comes close to anticipating this kind of response, but restricts the alternative possible explanation to justification based on observation. This leads him to another case, which he apparently takes to refute the contemplated alternative:

If I observe the truth of \underline{e} and then forget all about it, my evidence no longer includes \underline{e} . It is hard to see how evidence could discriminate between hypotheses in the way we want it to if it did not have to be known. (201)

The thinking here appears to be that since I did observe the truth of \underline{e} , the alternative theory has no way to account for my evidence not including \underline{e} at the present time. But our theory, $E = NPJ$, accounts for the case quite easily. If I have forgotten about \underline{e} (and am not currently observing \underline{e}), then I am no longer non-inferentially propositionally justified in believing \underline{e} . Hence, at the present time, \underline{e} is not evidence for me.

Williamson's next argument runs as follows:

If evidence required only justified true belief, or some other good cognitive status short of knowledge, then a critical mass of evidence could set off a kind of chain reaction. Our known evidence justifies belief in various true hypotheses; they would count as evidence too, so this larger evidence set would justify belief in still more true hypotheses, which would in turn count as further evidence.... The result would be very different from our present conception of evidence. (201)

Although the threat of a chain reaction may apply to a JTB account of evidence, it doesn't apply to the more restrictive account of evidence I am offering, which says that evidence propositions are non-inferentially justified propositions. No chain reaction is plausible under the NPJ account because the chain would require inferential justification, which is precluded from the start. Second, however, I am puzzled that Williamson presents this argument against the JTB approach to evidence, because the contemplated chain reaction seems almost as threatening under the E = K construal. Clearly, E = K does not restrict evidence to non-inferential knowledge. So if one's 'basic' evidence justifies belief in various true hypotheses that are also known, a very similar chain reaction looms. It would probably be a less extensive chain reaction, because some of the justified true beliefs might not satisfy the requirement of safety, and hence not amount to knowledge. But failure of the safety requirement might not trim the chain reaction very much. So a comparably worrisome scenario seems applicable to the E = K thesis.

Williamson next proceeds to defend the truth requirement on evidence:

That propositional evidence is knowledge entails that propositional evidence is true. That is intuitively plausible; if one's evidence included falsehoods, it would rule out some truths, by being inconsistent with them. One's evidence may make some truths improbable, but it should not exclude any outright. Although we may treat false propositions as evidence, it does not follow that they are evidence. No true proposition is inconsistent with my evidence, although I may think that it is. (201)

Is it correct that evidence should never exclude a truth? Everybody agrees that there can be misleading evidence. Why can't some evidence be deductively, as contrasted with inductively, misleading? The anticipated response is that since the probability of an item of evidence on the total evidence is 1, the negation of an item of evidence will have a probability of 0, and that probability will be unrevisable via conditionalization. But it won't be unrevisable full stop unless evidence stays fixed, and it is an open question whether a proposition that is evidence at one time must be evidence forever. Indeed, Williamson himself rejects the evidence permanence thesis, which he calls "monotonicity" (206, 218). If e is ostensibly observed at one time, it may qualify as evidence at that time. But at a later time, it may cease to qualify as evidence (and not merely through conditionalization). So allowing falsehoods as evidence doesn't entail that some truths are permanently consigned to epistemic oblivion.

A further reason Williamson gives for the view that all evidence consists of true propositions is that this explains the point of adjusting our beliefs to the evidence.

[A]djusting our beliefs to the evidence ... is a way of adjusting them to the truth. Although true evidence can still support false conclusions, it will tend to support truths. (202)

But $E = K$ isn't the only construal of evidence that can make sense of adjusting our beliefs to the evidence. $E = NPJ$ can make sense of it too. If we assume that justified propositions are likely to be true (an assumption that conforms with many theories of justification), then adjusting our beliefs to the evidence is a way of adjusting them to what is usually or mostly true. Evidence so construed will sometimes support false conclusions, but this holds equally under $E = K$, as the foregoing passage concedes.

Next Williamson considers the objection that perceptual appearances can provide evidence even without belief, which implies the falsity of $E = K$.

But, a critic may suggest, ... [in cases of perceptual evidence] my evidence includes e because it is perceptually apparent to me that things are that way, whether or not I believe that they are that way. Even if I do believe e, my evidence included e even before I came to believe it; according to the critic, I came to believe it because it was perceptually apparent.... We can ask the critic whether, for my evidence to include e, I must at least be in a position to know e? If so, then the critic's view does not differ radically from $E = K$. Given $E = K$, the evidence in my actual possession consists of the propositions which I know, but there is also the evidence in my potential possession, consisting of the propositions which I am in a position to know. The critic takes my evidence to be the evidence in my potential possession, not just the evidence in my actual possession. (202-203)

The critic imagined here endorses an approach not dissimilar to the one I am proposing. However, neither this critic nor I need accept Williamson's characterization of the approach. For example, it isn't being in a position to know e that makes it evidence; nor is having perceptual evidence a matter of the evidence merely being in my potential possession. Having evidence is a matter of (non-inferential) justification rather than knowledge. And in the perceptual case one is actually propositionally justified in believing e, not merely potentially so.

But Williamson poses a problem for this kind of view.

[S]uppose that I am in a position to know any one of the propositions p_1, \dots, p_n without being in a position to know all of them; there is a limit to how many things I can attend to at once. Suppose that in fact I know p_1 and do not know p_2, \dots, p_n . According to $E = K$, my evidence includes only p_1 ; according to the critic, it includes p_1, \dots, p_n . Let q be a proposition which is highly probable given p_1, \dots, p_n together, but highly improbable given any proper subset of them; the rest of my evidence is irrelevant to q . According to $E = K$, q is highly improbable on my evidence. According to the critic, q is highly probable on my evidence. $E = K$

gives the more plausible verdict, because the high probability of q depends on an evidence set to which as a whole I have no access. (203)

This argument says that the critic's view is implausible because the probability it assigns is based on a putative evidence set to which, as a whole, I (the cognitive agent) lack access. This is a highly unexpected argument to come from Williamson. It is a core part of his book to deny that either knowledge or evidence is transparent. As he writes in the chapter on skepticism, "Rational thinkers are not always in a position to know what their evidence is; they are not always in a position to know what rationality requires of them...." (164). So how can he regard it as a deficiency in the critic's position -- or the closely related $E = NPJ$ position -- that one might not have access to one's evidence as a whole? It could still be one's evidence even if one doesn't have access to all of it.

Let me offer an example of my own to defend the greater plausibility of $E = NPJ$ as compared with $E = K$. Consider someone who is quite cautious in doxastic matters, a person I shall call a diffident doxastic agent (DDA). One such person suspends judgment about everything. We needn't make our case, however, by reference to such an atypical, psychologically unrealistic, individual. Consider instead a DDA who has only mild states of credence, none of which rise to the level of belief (wherever that level of credence is set). It follows from the $E = K$ thesis that such a DDA has no evidence at all, because knowledge requires belief.

Is it plausible to hold of such a DDA that he totally lacks evidence? This DDA, we may assume, has a full panoply of perceptual experiences. These experiences, I submit, provide him with evidence for many propositions. It might be debated exactly which propositions they are evidence for -- physical-object propositions, internal state propositions, etc. But a resolution of this debate doesn't affect the present argument. If, in virtue of these perceptual experiences, he is non-inferentially justified in believing any propositions, then according to $E = NPJ$, these propositions are items of evidence for him. This is a much more plausible view of his evidential state than $E = K$ provides.

I conclude that Williamson does not offer a compelling rationale for the $E = K$ thesis. A preferable view, notably different from $E = K$, has been outlined, which accounts for all the intuitive "data" Williamson adduces for $E = K$. No doubt there are many other possible views as well, intermediate between $E = K$ and $E = NPJ$. The main point, however, is that no convincing argument has been given for $E = K$.

7. Evidence and Skepticism

One of the chief attractions of $E = K$ in Williamson's eyes is the help it offers in dealing with skepticism. A standard form of skepticism contends that one cannot know things about the external world even in the Good Case (where things appear generally as they ordinarily do, and are that way) because one's evidence in the Good Case is identical to one's evidence in the Bad Case (where things appear as they ordinarily do, but are some other way). So how can one tell that one is in the Good Case rather than the Bad Case? The $E = K$ thesis makes it easy to reject the alleged evidential symmetry of the

Good Case and the Bad Case, thereby undercutting the skeptic's argument. The skeptic is not entitled to assume that we have no more evidence in ordinary cases than in their skeptical counterparts, because the anti-skeptic contends that we do have more knowledge – hence more evidence – in ordinary cases than in their skeptical counterparts.

If our critique of the reasons given for $E = K$ stands, $E = K$ no longer looks so promising, and that seems to offer comfort to the skeptic. Is that the stance I mean to encourage? First, I don't mean to take a firm stance here on the best way(s) to reply to skepticism. Second, the alternative construal of evidence I have placed on the table may also leave room for the same reply to the skeptic Williamson favors, viz., that one has more evidence in the Good Case than in the Bad Case. Let's see why this is so.

Given $E = NPJ$, the extent of one's evidence in the Good Case and the Bad Case depends on the chosen account of non-inferential propositional justification. On certain approaches to justification, one would expect there to be more non-inferential propositional justification in the Good Case than in the Bad Case.¹⁵ In particular, this is likely to hold for certain forms of externalism about justification, especially a form of reliabilism in which reliability is fixed by truth-ratios of belief in the world in question. An ordinary-world agent is likely to have more non-inferential propositional justification vis-à-vis perceptual propositions (about physical objects) than an experientially similar envatted agent would have.¹⁶ Thus, the asymmetry of evidence between the Good Case and the Bad Case is defensible under some elaborations of the $E = NPJ$ interpretation of evidence. Williamson, of course, would not be entirely happy with this upshot. He likes the $E = K$ thesis in part because it fits with his "knowledge first" approach to epistemology in general. The $E = NPJ$ interpretation of evidence does not fit with the "knowledge first" credo. I don't find that a liability, because I am unpersuaded of the "knowledge first" theme. But it's a major virtue of Williamson's approach that it puts all these interesting theses on the table, and offers highly intriguing arguments for them.¹⁷

References

- Goldman, A. (1976). Discrimination and Perceptual Knowledge. Journal of Philosophy 73: 771-791.
- Goldman, A. (1979). What Is Justified Belief? In G. Pappas, ed., Justification and Knowledge. Dordrecht: Kluwer.
- Goldman, A. (1986). Epistemology and Cognition. Cambridge, MA: Harvard University Press.
- Goldman, A. (1992). Liaisons: Philosophy Meets the Cognitive and Social Sciences. Cambridge, MA: MIT Press.
- Goldman, A. (forthcoming). Immediate Justification and Process Reliabilism. In Q. Smith, ed., Epistemology: New Essays. New York: Oxford University Press.
- Neta, R. and Rohrbaugh, G. (2004). Luminosity and the Safety of Knowledge. Pacific Philosophical Quarterly 85: 396-406.
- Nozick, R. (1981). Philosophical Explanations. Cambridge, MA: Harvard University Press.
- Roush, S. (2005). Tracking Truth. Oxford: Oxford University Press.
- Silins, N. (2005). Deception and Evidence. In J. Hawthorne, ed., Philosophical Perspectives, 19, Epistemology (pp. 375-404). Malden, MA: Blackwell.
- Whitcomb, D. (2005). Factivity without Safety. Unpublished manuscript, Rutgers University Department of Philosophy.
- Williamson, T. (2000). Knowledge and Its Limits. Oxford: Oxford University Press.

¹ All page numbers refer to this work.

² Williamson would say that the JTB account omits a safety condition, which he regards as a non-circular necessary condition.

³ Thanks to Dennis Whitcomb for emphasizing the need to be clear on this point.

⁴ Thanks here to Karson Kovakovich.

⁵ Igal Kwart offered this argument.

⁶ Most regrettably, Ginet was not credited in the original publication (Goldman, 1976). I belatedly tried to set the record straight by adding an endnote with the credit in a subsequent reprinting of “Discrimination of Perceptual Knowledge” (Goldman, 1992). Few people, however, seem to have noticed that endnote. Others, meanwhile, have quite properly called attention to the fact that credit for the example should go to Ginet.

⁷ A somewhat similar argument appears in Whitcomb (2005).

⁸ I emphasize the relevance to knowledge of both global reliability and local (modal) reliability in (Goldman 1986).

⁹ Also at one point Williamson seems to confuse process reliabilism with the safety brand of reliabilism. He writes: “If one believes p truly in a case α , in which other cases must one avoid false belief in order to count as reliable enough to know p in α ? There is no obvious way to specify in independent terms which other cases are relevant. This is sometimes known as the generality problem for reliabilism” (100). This isn’t exactly right. The generality problem is a problem for process reliabilism specifically, not safety theory. It is the question of which process (or method) type is the correct type to use when fixing the reliability of a belief’s generating process. That Williamson gets this terminology wrong – and otherwise neglects process reliabilism – suggests that he has little interest in this theoretical angle.

¹⁰ In particular, it would have strong conditional reliability properties. The conditional reliability of an inference procedure is the proportion of times its belief outputs are true given that all the inputs to it are true. See Goldman (1979).

¹¹ Of course, introducing a global reliability constraint poses the well-known generality problem. I have nothing new to offer here on that topic.

¹² There are well-known counterexamples to the necessity of a simple reliable-process condition. So a subtler form of reliable-process condition is required. We needn’t speculate here as to what form such a condition should take. It suffices to point out by example that the safety condition isn’t adequate, precisely because it doesn’t preclude inferior processes being used in the actual situation.

¹³ My competing construal of evidence does not mesh with Williamson’s evidentialist project of using the notion of evidence in a theory of justification, because my competing construal of evidence runs in the opposite direction: it uses justification in the account of evidence. But if Williamson is to persuade us of his evidentialist project, he needs arguments to support it. Section 9.6 provides such an argument, offering a challenge to find an alternative account of evidence that fits as well with the “data” adduced there as $E = K$ does. I claim that my competing construal of evidence meets this challenge, thereby undercutting $E = K$ and, indirectly, the evidentialist program.

¹⁴ I set aside the question of how such an approach to evidence would deal with cases in which prima facie justifiedness is defeated.

¹⁵ Silins (2005) points out that there are various weaker versions of evidential externalism than the one Williamson offers. For example, an evidential externalist might accept the thesis that one's propositional evidence is what one justifiedly and truly believes (2005: 378). What is crucial to evidential externalism is that what evidence one has is sensitive to the environment one is in. Obviously, this isn't equivalent to the E = K thesis. Silins himself weakly favors evidential internalism.

¹⁶ For a defense of a process reliabilist account of immediate (non-inferential) justification, see Goldman (forthcoming). However, that defense does not address the question of how reliability should be fixed. So it is silent on, and has no unambiguous consequences for, the present issue of skepticism.

¹⁷ Special thanks to Dennis Whitcomb and Karson Kovakovich for extensive and astute comments on one version of this paper, to Alex Jackson, Igal Kvat, Ernest Sosa, and other participants in the Spring 2006 Epistemology seminar at Rutgers University, and, finally, to Duncan Pritchard for valuable comments on the penultimate draft.