

# 9

## Do We Have Reason to Doubt the Importance of the Distinction between A Priori and A Posteriori Knowledge?

A Reply to Williamson<sup>1</sup>

*Paul Boghossian*

1

### *The resilience of the a priori*

The category of a priori justification has proven remarkably resilient.<sup>2</sup> Despite persistent and powerful attacks on its possibility, reality, and even its very coherence, it has remained part of the basic conceptual framework of philosophy. What explains this remarkable staying power?<sup>3</sup>

One of the reasons why the category of a priori justification has proven so resilient derives from the pull of certain intuitive *examples*: when we contemplate the pair of propositions

(i) It's sunny

and

(ii) If it's sunny, then it's sunny,

we can't help but be struck by the thought that experience plays a very different role in the way we typically know these two propositions. Roughly speaking, in the case of the first, we need to open our eyes to see whether it is sunny; in the case of the second, we don't.<sup>4</sup>

Furthermore, these examples are projectible: on their basis, we know how to classify *further* cases as a priori or a posteriori, and we achieve a high level of non-collusive convergence on those further classifications. For example, we are fairly confident that

(iii) All squares have four sides.

- (iv) Nothing can be red and green all over at the same time.
- (v) In a Trolley Scenario, it is morally wrong to throw the heavy man off the bridge in order to save the five innocents trapped below.
- (vi) Two circles of different radii can at most intersect at two points.
- (vii) It is necessary that whoever knows *p* believes *p*.

all fall on the a priori side, while

- (viii) All wooden tables are squares.
- (ix) This volume of *Who's Who* is red.
- (x) In all societies, it is considered morally wrong to throw the heavy man off the bridge in a Trolley Scenario.
- (xi) There are two intersecting circles on this page.
- (xii) Whoever knows *p* enjoys knowing *p*.

fall on the a posteriori side.

While we have a fairly good handle on *how* we know propositions in an a posteriori way, we are not entirely sure how we know propositions a priori: How, for example, do we know (ii) without experience? Do we examine our concept of 'if'? Do we have an intuition that all instances of the sentence 'If *p*, then *p*' are necessarily true? Or is the story something else altogether?

To concede this is to say that we are not sure how to *explain* our capacity for experience-independent justified belief. But the datum itself, that there are such beliefs, seems firm enough.

All of this gives us a very strong basis on which to be confident that, through these examples, we have got hold of a genuine distinction, with the only question being, How exactly should that distinction be theoretically formulated and how should its various instances be explained?

And that is why, despite all the impressive and by no means fully answered challenges to the notion of the a priori, it continues to have a strong grip on us.

### *A new style of skepticism*

Timothy Williamson has developed a novel form of skepticism about the a priori, one that sees no need to deny any of the observations made above—indeed, he makes most of them himself.

According to Williamson, although the robustness and projectibility of the examples warrant our confidence that we have got hold of a genuine distinction, they do not warrant our believing that the distinction is a significant one, of any great importance to epistemology.

Indeed, Williamson claims to be in possession of an argument that shows that *whatever* distinction we have got hold of through our examples it must be a *shallow* one, one that serves to obscure important epistemological generalizations, rather than illuminate them.

This is a genuinely novel form of skepticism about the a priori—a posteriori distinction: it accommodates our strong conviction that we have got hold of a real distinction, even while it purports to undermine our sense of the importance of that distinction.

In this paper, my main concern will be to show that Williamson's central argument for his novel form of skepticism about the a priori doesn't succeed.

However, en route to presenting his central argument, Williamson develops an interesting objection of another, more old-fashioned type, one that is directed against the very possibility of formulating a satisfactory theoretical (top-down) distinction between the a priori and the a posteriori. (As we shall see in a moment, this more traditional argument turns on the classic difficulty in saying what notion of experience is involved in the core apriorist notion of 'a justification that is independent of experience.')

Since this argument of Williamson's is of strong independent interest, and since it will enable us to introduce some basic concepts that will be important later on, I will begin by looking at it, before turning to his central argument.

## 2

### *A problem for characterizing 'experience'*

The fundamental notion of a priori justification is one that applies to a *way of knowing* or *justifiably believing* a proposition. According to the traditional explication,

S's belief that p is justified a priori if and only if S's belief that p is *justified in a way that is independent of experience*.

Derivatively, this notion of apriority can be applied to a *proposition*:

The proposition that p is a priori if and only if p *can* be justifiably believed, by creatures like us, in an a priori way.

As is well known, even this rudimentary characterization immediately faces a glaring issue. According to our intuitive examples, typical ways of justifiably believing that if it's sunny then it's sunny are a priori. But my justifiably believing this proposition will certainly depend on experience in *some* sense: after all, my very ability to *entertain* the proposition that if it's sunny then it's sunny (plausibly) depends on my having had some experience with the sun.

The standard way of dealing with this issue is to say that my knowledge that if it's sunny then it's sunny can still count as a priori, even if experience is playing some role in my arriving at it, provided that it is playing at most an *enabling* role, but not a *justificatory* or *epistemic* role.<sup>5</sup>

What does it mean for experience to play a merely enabling role? Here there are two options that are sometimes conflated.

On the one hand, we might have a purely *negative* characterization of 'enabling' in mind, one that applies to *any* role that experience may play in generating belief apart from providing a justification for the belief in question.

On the other hand, we might have in mind a characterization that is in part *positive*, for example, having to do with the role that experience plays in enabling one to grasp and therefore entertain certain propositions.

On the negative view, *any* role that's non-epistemic is enabling. On the positive view, by contrast, it would be an open question whether experience occasionally plays a role that is neither epistemic nor enabling.

Williamson assumes that the friend of the a priori is committed to a *positive* characterization of 'enabling': one of his big themes, as we shall see, is that experience often plays a role that is *intermediate* between an enabling role and an epistemic role, a possibility that couldn't arise on the negative characterization.

However, Williamson doesn't explain why the friend of the a priori couldn't be content with simply insisting that experience *not* play an *epistemic role*, without also giving a positive characterization of all the roles that are thereby left over. Of course, Williamson's opponent would owe us an account of what it is for experience to have an 'epistemic role', but that is a perfectly general burden that any epistemologist must bear.

If we ask which positive characterization of 'enabling' Williamson assumes the friend of the a priori must have, the answer seems to be (although he doesn't quite spell it out) that it is a broad notion of *access* to a proposition.

Since it won't matter for present purposes, let us grant Williamson this positive, access-involving notion of an enabling role for experience. Let us turn instead to asking what the word 'experience' should be taken to cover in the expression 'experience that does not play an epistemic role'. Williamson says the following (it will be useful to quote him in full):

Another issue is how widely to apply the term 'experience'. It is mainly associated with 'outer' experience, involving perception by the usual five senses, but why should it exclude 'inner' experience, involving introspection and reflection? After all, one's knowledge that one is in pain is presumably a posteriori, even though the experience on which it depends is inner. Excluding inner experience by stipulation, without reference to any deeper epistemological difference, is liable to make the distinction between a priori and the a posteriori knowledge epistemologically superficial. Inner and outer experience will therefore provisionally be treated on an equal footing.

One might worry that if inner experience is included, our experience of reflecting on the proposition that if it is sunny then it is sunny will play an evidential role in our knowledge that if it is sunny then it is sunny, and that Mary's experience of calculating that  $289 + 365 = 654$ , on paper or in her head, will play an evidential role in her knowledge that  $289 + 365 = 654$ . Presumably, the response is that the role is purely enabling. The relevant evidence is not the psychological process of reflecting or calculating, but rather in some sense the non-psychological logical or mathematical facts to which that process enables one to have access.

On further thought, however, that response causes more problems than it solves. For what prevents it from generalizing to outer experience? For example, part of the evidence that a massive comet or asteroid collided with the Earth about 250 million years ago is said to be that certain sediment samples from China and Japan contain certain clusters of carbon atoms. That those samples contained those clusters of atoms is a non-psychological fact. Of course, in some sense scientists' outer experience played a role in their access to the fact. But, by analogy with the logical and mathematical cases, the relevant evidence is not the psychological process of undergoing those outer experiences, but rather the non-psychological physical facts to which that process enables us to have access. The role of the outer experience is purely enabling, not evidential. If so, what would usually be regarded as paradigm cases of a posteriori knowledge risk reclassification as a priori.

The threat is not confined to theoretical knowledge in the natural sciences. Even for everyday observational knowledge, it is a highly controversial move to put the psychological process of undergoing the outer experience into the content of the perceptual evidence we thereby gain. What we observe is typically a non-psychological fact about our external environment, not a psychological fact about ourselves.<sup>6</sup>

Williamson's argument may be rendered as follows:

1. In the locution 'justification that is epistemically independent of experience', 'experience' might apply just to 'outer experience', or it might apply both to 'outer experience' and to 'inner experience.'
2. If it applies only to 'outer experience'—thereby including perception by the usual five senses, but excluding introspection, reflection, and mental calculation—then the distinction between the a priori and the a posteriori will be superficial, since we will have to classify both our knowledge of mathematics and our knowledge that we are in pain (when we are) as a priori.
3. If it applies both to 'outer' and to 'inner' experience, then we will have to say that the role of the mental process of mathematical *calculation* is merely enabling, giving us access to the relevant mathematical facts in a computation, rather than playing a role in the justification of the result that we arrive at.
4. However, if we take that line about mental calculation, we will have to generalize it to the case of outer experience as well, with the result being that what were supposed to be paradigm cases of a posteriori knowledge (for example, that an asteroid collided with the Earth 250 million years ago) end up being reclassified as a priori.

Williamson's argument poses a dilemma for any theoretical, top-down characterization of the notion of 'justification that is epistemically independent of experience'. The dilemma turns on whether 'experience', in this phrase, is to be construed so as to include or exclude 'inner' experience. Either way, the thought is, it appears impossible to get a characterization of the distinction that lines up with the distinction indicated bottom-up by our intuitive examples.

#### *Discussion of Williamson's problem for characterizing 'experience'*

A central difficulty with this argument is that it seems to assume that there is some natural, joint-carving, distinction between 'inner' and 'outer' experience, one which forces us to choose, en bloc, whether to include or exclude 'inner' experience. But there is no such distinction.

In fairness to Williamson, it should be admitted that the existence of such a distinction has often been presupposed by writers in the philosophy of mind (including the present author). On reflection, though, it is hard to see what it could consist in.

For what exactly is the distinction between 'inner' and 'outer' experience supposed to track? Williamson doesn't say, but the only candidate would appear to be the status of the *intentional objects* of the relevant types of experience. Thus, visual perception would count as 'outer' because its intentional objects are outside the head, whereas introspection would count as 'inner' because its intentional objects, such as pains, are inside the head.

But the distinction, as so characterized, is not well defined, and doesn't correspond well to the use that Williamson makes of it. Numbers, for example, on most views, including Williamson's own, are outside the head—but he regards mental calculation, whose intentional objects are numbers, as 'inner'.

And where should we say that memory falls on this way of drawing the distinction between inner and outer? Some of the intentional objects of memory are 'outer', when they concern worldly events, and some are 'inner', when they concern past mental states. So, memory could not be classified neatly as either inner or outer.

However, if the distinction between inner and outer experiences is not based upon the status of their characteristic intentional objects, then I don't see what the distinction could be based upon. In the only other sense that's available, all experiences are inner, in that they are inner states of the mind of the subject.

Williamson might object that it is not *he* who is endorsing the 'inner'–'outer' distinction, but rather the theorist of the a priori; and he could point to his use of shudder quotes around 'inner' and 'outer' to support his claim. But this response requires him to explain why he thinks *anyone* is committed to the distinction, and in particular why he thinks the friend of the a priori is.

Another possible response on Williamson's behalf would have it that the distinction between inner and outer is given not by a principle, but by a list: sense experience is outer and all the rest—including introspection, memory, intuition, imagination, and calculation—are inner.

I am not averse to specifying distinctions via lists, as I will explain further below. But if we did so in this case, it would be open to the friend of the a priori to insist that while 'experience' will cover at least some 'inner' experiences, as given by the list, it won't cover them all. So, while it will include substantive uses of memory and introspection, it *won't* include intuition, most uses of the imagination, and most uses of the capacity for mental calculation: those faculties, it is natural to think, provide a priori warrants for belief, when they provide any warrant at all.

Wouldn't such a response count as hopelessly gerrymandered?

Well, it would count as a gerrymandered response if the stipulated distinction between inner and outer experience had been a natural one, one that carved experience at its natural seams. But, as we have seen, it is no such thing.

There remains a good question, of course, whether we can articulate a *principle* that will capture the distinction between those types of experience that are included in 'epistemically independent of experience', or whether we have to rest content with specifying the distinction via a list.

While there have been some promising attempts, it would be fair to say that, as of yet, no one has formulated a principle that captures the intuitive distinction between a priori and a posteriori sources of warrant correctly.<sup>7</sup> But that is a weak reason to think it can't be done, or that the distinction at issue is not a good one. And it is worth bearing in mind that, as Allan Gibbard has wisely reminded me, there isn't in the end such a large difference between a principle and a list, since a principle is just a list of one.<sup>8</sup>

### *Justifiers as propositions*

Now, while I believe that this particular argument of Williamson's can be answered, it is worth pausing for a moment on an important lesson that is implicit in it.

Williamson's argument identifies a crucial danger for the friend of the a priori, one that perhaps has not been previously sufficiently emphasized. It arises at the point at which the friend seems driven to say that, in the case of a mental calculation, the *justifier* for the belief that the sum of 289 and 365 is 684 is not the *mental process* of calculating those numbers, but rather just the relevant arithmetical propositions to which the calculation provides access.

And the important lesson is that once we begin to think of the justifiers of our beliefs as propositions rather than as mental states, the threat of overgeneralizing a priori belief becomes hard to avert.

For, as we saw above, if the justifier for my 'empirical' belief that an asteroid collided with Earth is just the proposition that samples of sediment contain certain clusters of carbon atoms, rather than any visual observations that justified the premise belief in the first place, then we would have to say that, even in 'empirical' knowledge, sensory experience itself plays at most an enabling 'access' role, and we would thereby have lost any interesting contrast between a priori justification and a posteriori justification.

A priori belief is belief that is held with a priori justification. The relevant notion of justification here can't be merely that of a proposition, but must include reference to the mental states that have those propositions as intentional objects. If we were restricted to making reference only to propositions, the natural thought would be that an a priori justified belief would be a belief that is justified solely by an a priori justified proposition. But which propositions would those be? The resultant view would be hopelessly impredicative.

This is an interesting result, one that, so far as I know, neither friend nor foe of the a priori has sufficiently emphasized. The friend of the a priori cannot afford to adopt the view that justifiers are propositions, rather than mental states, for that commitment is incompatible with an interesting distinction between a priori and a posteriori justification.

In the argument as Williamson presents it, he suggests that the case of mathematical calculation will drive the friend of the a priori to think of justifiers as propositions. As we have seen, his argument for this is based on the optional demand that the friend make up her mind on whether all 'inner' experience is included in the term 'experience', and so we were able to skirt around it.

These days, though, the view that justifiers are propositions, rather than mental states, is a view that has very broad appeal to epistemologists for reasons having nothing to do with the debate about the a priori.<sup>9</sup>

I myself do not see that appeal, and not merely because of my views on the a priori. To me it seems obvious that my justification for believing that, for example, there is a red wall in front of me, is irreducibly dependent on my having a *visual impression* of there being a red wall in front of me.

However, there is a case that has sometimes seemed, even to friends of the a priori, to provide strong motivation for the justifiers as propositions view. The case is still based on the example of mathematics, but stems not from reflecting on the distinction between inner and outer experience, but by looking at the role of memory and perception in proof.

In the remainder of this section, I will lay out this motivating argument and suggest a way of blocking its conclusion.

*Proof and memory*

In the case of some lengthy proofs, it is not possible for creatures like us to carry the whole thing out in our minds. Unable to keep all the steps of the proof in mind, we need to write them down and look them over. In such cases, memory of the earlier steps in the proof, and perceptual experience of what has been written down, enter into the full explanation of how we arrive at our belief in the conclusion of the proof.

And these facts raise a puzzle that has long worried theorists of the a priori: How could belief arrived at on the basis of this sort of lengthy proof be a priori warranted? Won't the essential role of memory and perceptual experience in the process of carrying out the proof undermine the conclusion's alleged a priori status?<sup>10</sup>

To acquiesce in a positive answer to this question would be counterintuitive. By bottom-up standards, lengthy proof can deliver an a priori warrant for its conclusion just as well as a short proof can. The theoretical puzzle is to explain how it can do so given the psychologically necessary role of perceptual and memory experiences in the process of proof.

One natural way of responding to this puzzle has been developed by Tyler Burge (1993) (Burge focuses on memory, but much the same could be said for perception):

Memory does not supply for the demonstration propositions about memory, the reasoner or past events. It supplies the propositions that serve as links in the demonstration itself. Or rather, it *preserves* them, together with their judgmental force, and makes them available for use at later times. Normally, the content of the knowledge of longer demonstration is not more about memory, the reasoner, or contingent events than that of a shorter demonstration. One does not justify the demonstration by appeals to memory. One justifies it by appeals to the steps and the inferential transitions of the demonstration... In a deduction, reasoning processes' working properly depends on memory's preserving the results of previous reasoning. But memory's preserving such results does not add to the justificational force of the reasoning. It is rather a background condition for the reasoning's success.<sup>11</sup>

What exactly is the distinction between preservative memory and substantive memory that Burge is invoking here? As he notes, it is reminiscent of the distinction, familiar from psychology textbooks, between episodic memory and semantic memory. We distinguish between remembering *being told* that the Battle of Hastings was in 1066, and remembering that the Battle of Hastings was in 1066. The former is substantive memory, which can import the event of being told about the battle into reasoning and can justify believing that it happened. The latter is just (preservative) memory of the proposition itself. It only serves to preserve the proposition itself (along with its 'judgmental force') and does not contribute any new justification.<sup>12</sup>

Given this distinction, we can say that the reason why a long proof is able to provide a priori justification for its conclusion is that the only use of memory that is essential in a long proof is preservative memory, rather than substantive memory. If substantive memory of the act of writing down a proposition were required to arrive at justified belief in the conclusion of the proof, that might compromise the conclusion's a priori status.

However, once we lean on the notion of preservative memory in this way, we might be tempted to think, as Burge clearly does, that the justifiers in a proof must be

just the mathematical propositions themselves, rather than any sort of accompanying mental state.

For while we think we understand how a *proposition* might be the object of a merely preservative memory, we might find ourselves mystified by how a *mental state* or *event* could be the object of such a memory. It looks as though remembering a mental state or event will involve a substantive use of memory, rather than just a preservative one.

If you are impressed with this line of thought, you might find yourself driven to think that, in the case of a mathematical proof, the justifiers must consist in propositions, rather than in any mental states, like those of intuition or mental calculation, which might have those propositions as their intentional objects.

And, as I emphasized earlier, once you have arrived at this point, it is bound to look as though Williamson's difficulty simply reasserts itself. For, if the correct view of justifiers in the mathematical case is that they are propositions, that would presumably also be the correct view in the empirical case. However, if one adopts such a view in the empirical case, the role of perceptual experience, just like that of preservative memory, would have to be thought of as simply *providing access* to the relevant justifiers (the propositions themselves), rather than playing any essential role in justifying them. And at that point we would be landed back again with the difficulty of dramatically over-generating the beliefs that would have to count as a priori justified.

### *Resolving the puzzle*

The correct reply to this argument is to reject its assumption that the only things that preservative memory is capable of preserving are propositions (along with, at most, their judgmental force).

Rather, there is no reason why we should not regard preservative memory as being able to preserve a proposition not only with its judgmental force but *also* along with its *epistemic status*. Doing so will allow us to say that justifiers are mental states, but that no substantive memory of those states is required in order for the proposition that is justified by them to be later used in warrant-transmitting reasoning. Once the proposition is justified by the relevant mental states, it, along with its judgmental force and epistemic status, is preserved in memory for use in later reason-generating inference.

This way of construing matters is reinforced by the consideration that even at the point of *onset* of belief in the proposition, the justification of the proposition by the relevant mental state doesn't depend on a second-order *awareness* of that mental state, but only on the existence of the first-order mental state itself.

To illustrate the claims I am making here, consider an empirical case involving justification by visual perception. Suppose you believe, as I do, that my justification for believing that these sediment samples' containing certain clusters of carbon atoms can be irreducibly dependent on my having a certain sort of visual experience of them (I will assume for the purposes of argument that this proposition is purely observational). When I later come to use this proposition in an inference to the best explanation, to conclude that the Earth collided with an asteroid 250 million years ago, it would be counterintuitive to say that substantive memory

of my having had the earlier visual experiences enters into the justification for my conclusion. Rather, the original experiences serve to justify the observational premise, and then that positive epistemic status gets simply preserved in mounting the inference on its basis. My justification for believing the conclusion rests on the justification that I have for the premises, which in turn depends on my prior visual experiences; the justification doesn't rest on my remembering those experiences or, at the point of origin, on my being aware of having those experiences (even if it may not be possible to have an experience without having some awareness of it).

So, even in a humdrum empirical case, it is natural to think that preservative memory can preserve propositions not merely with their judgmental force but also with their epistemic status intact.<sup>13</sup>

And, of course, what goes for empirical inference also goes for mathematical proof. Just because a mental state was involved in the justification of an earlier step in the proof doesn't mean that one needs to have a substantive memory of that mental state in order to justifiably infer from that premise to a later step in the proof. It is sufficient that the premise was justified, and that its justificational status was preserved by preservative memory.<sup>14</sup>

Although I have rejected Williamson's argument against the possibility of formulating a cogent theoretical distinction between the a priori and the a posteriori, it has provided us with an opportunity to learn an important lesson: Taking the distinction between the a priori and the a posteriori seriously is not compatible with an increasingly influential conception of justifiers as propositions. Rather, taking the distinction seriously depends on a traditional view of justification as stemming from mental states and events, some types of which are taken to be sources of a posteriori warrant, while others of a priori warrant.

The discussion itself illustrates a more general phenomenon to which we will have occasion to return: that the distinction between the a priori and the a posteriori interacts very closely with issues in other areas of epistemology, and in philosophy more generally. Commitments arrived at elsewhere, on grounds having nothing to do with the a priori, may put pressure on one's ability to make sense of the notion. Once the dependence is made clear, a useful discussion about the trade-offs involved can ensue.<sup>15</sup>

### 3

#### *Williamson's central argument*

Having attempted to show that no top-down characterization of the a priori–a posteriori distinction can hope to succeed in lining up with the paradigm intuitive examples, Williamson changes tack. He next attempts to argue that whatever distinction is marked by the paradigm examples, it cannot be of epistemological significance. How does Williamson propose to support this strong and surprising conclusion?

Williamson's argumentative strategy is striking: it is to describe a pair of examples with the following features:

- (A) One of them is a clear case of someone's justifiably believing a proposition in an a priori way;
- (B) The other is a clear case of someone's justifiably believing another proposition in an a posteriori way;
- (C) The ways of justifiably believing involved in the examples are *broadly representative* of the ways in which propositions are justifiably believed a priori and a posteriori, respectively; and, finally,
- (D) There is no epistemologically interesting difference between the two ways of believing specified.

I am prepared to concede that, if we really did have a pair of examples with all these features, we would have to admit that the distinction between a priori and a posteriori ways of knowing is not epistemologically interesting.

I will argue, however, that Williamson's pair of examples lack *any* of these features. Williamson's examples involve the following two truths:

- (1) All crimson things are red.
- (2) All recent volumes of *Who's Who* are red.

He describes the case of Norman who is said to come to know (1) in the following distinctive way: Having acquired competence with the concepts involved in (1) independently of one another, he comes to know (1) by using his imagination in the following way: he imagines a crimson thing, notes that it is red, under the imaginative supposition, and so concludes that (1) is true.

Williamson is aware that a reader will worry that Norman may have imagined a peripheral shade of crimson. If one shade of crimson is red, it does not follow that all are. In reply, he says,

The relevant cognitive skills must be taken to include sensitivity to such matters.

Williamson then wants to give a very similar story for Norman's knowledge of (2).

Norman is as already described. He learns the complex phrase 'recent volumes of *Who's Who*' by learning 'recent', 'volume', '*Who's Who*' and so on. He is not taught any rule like (2), connecting 'recent volume of *Who's Who*' and 'red'. Through practice and feedback, he becomes very skilful in judging by eye whether something is a recent volume of *Who's Who* (by reading the title), and whether something is red. Now Norman is asked whether (2) holds. He has not previously considered any such question. Nevertheless, he can quite easily come to know (2), without looking at any recent volumes of *Who's Who* to check whether they are red, or even remembering any recent volumes of *Who's Who* to check whether they were red, or any other new exercise of perception or memory. Rather, he assents to (2) after brief reflection along something like the following lines. First, Norman uses his skill in making visual judgments with 'recent volume of *Who's Who*' to visually imagine a recent volume of *Who's Who*. Then he uses his skill in making visual judgments with 'red' to judge, within the imaginative supposition, 'It is red'. This involves the same general human capacity as before to transpose 'online' cognitive skills originally developed in perception into corresponding 'offline' cognitive skills subsequently applied in imagination. No episodic memories of prior experiences, for example of recent volumes of *Who's Who*, play any role. As a result of the process, Norman accepts (2). Since his performance was sufficiently skilful, background conditions were normal, and so on, he thereby comes to know (2). (2013: 296)

And so Williamson concludes:

The problem is obvious. As characterized above, the cognitive processes underlying Norman's clearly a priori knowledge of (1) and his clearly a posteriori knowledge of (2) are almost exactly similar. If so, how can there be a deep epistemological difference between them? But if there is none, then the a priori–a posteriori distinction is epistemologically shallow. (2013: 296–7)

Williamson goes on to give a theoretical *diagnosis* of what makes these two cases of putative knowledge similar and of why they pose a problem for the traditional distinction between the a priori and the a posteriori. He contends that it is because, in both of them, experience plays a role that is *intermediate* between a purely enabling role and an epistemic one. I will come back to this diagnosis.

### *Discussion of Williamson's central argument*

Williamson's argument is interesting in a number of ways, especially in its focus on the imagination as a source of a posteriori justification. While the imagination has always been regarded as a potential source of a priori justification, much less thought has been given to the way in which it can sometimes serve as a source of a posteriori justification.

While I agree with Williamson that the imagination can serve as a source of a posteriori justification, I don't believe that its use is anywhere near as ubiquitous as his argument requires.

Nor do I believe that, when it does serve as a source of such justification, it functions in the basically externalist, Reliabilist way in which Williamson's stories about Norman presuppose.

I will start by looking critically at (B) and move on to show that none of propositions (A) through (D) is satisfied.

In my experience, many philosophers have greeted with raised eyebrows Williamson's claim that Norman could come to justifiably believe (2) in the way described. That is my own reaction as well.

Norman is said to *imagine* a single instance of a volume of *Who's Who*, note that, within the imaginative supposition, its color is red, and conclude on this basis that *all* recent volumes of *Who's Who* are red. And his belief is supposed to be justified, indeed, justified with a strength sufficient for knowledge.

That should certainly seem *prima facie* puzzling. We wouldn't credit Norman with justified belief that *all* recent volumes of *Who's Who* are red on the basis of his vivid and accurate *perception* of a *single* red volume. Why would we be so much more epistemically generous with regard to his act of *imagining* a single such volume? Isn't imagining even less constrained by actual truth than perception?

In fairness to Williamson, we should emphasize that there is supposed to be more to Norman's method of forming a belief in (2) than a single act of imagination. On those bare bones, Williamson sprinkles a heap of reliability-securing counterfactuals. He is not explicit about what they are, but perhaps something along these lines: If all the recent volumes of *Who's Who* had not been red, but had rather been green, Norman would not have imagined a red volume of *Who's Who*, but rather a green one; and if they not been of any uniform color, then Norman would have drawn a blank while trying to imagine one. And so on.

According to Williamson, then, while Norman's single act of imagining may not, by itself, be enough for Norman to justifiably believe (2), Norman's act of imagining, along with the truth of these counterfactuals, *is*.

To my mind, and as I will explain in much greater detail below, this verdict looks less like the deliverance of intuition than the deliverance of a prior allegiance to the dubious doctrine of Epistemic Reliabilism—roughly, the view that the reliability of a belief-forming method, regardless of whether the agent is aware of its reliability, is *sufficient* for a belief based upon that method to be epistemically justified.

Since I take there to be clear-cut and decisive counterexamples to Reliabilism—for example, Keith Lehrer's Mr. TrueTemp—I am not moved by Williamson's claims about Norman's reasonable belief in (2).<sup>16</sup>

Williamson's choice of name for his fictional subject is risky, since it can't help but bring to mind one of the most well-known and persuasive counterexamples to Reliabilism, namely, Lawrence Bonjour's clairvoyant Norman.<sup>17</sup> No doubt meant as an act of defiance, it nonetheless vividly reminds us of all the reasons we have for rejecting the epistemic upshot that Williamson claims for his examples.

I take the existence of such counterexamples to Reliabilism to derive from the fact that epistemic justification is a normative notion, one that is internally tied to notions of praise and blame. So, I don't agree with Williamson that Norman ends up reasonably believing (2).<sup>18</sup>

From this perspective, Williamson's sprinkling of Reliabilist counterfactuals onto Norman's single act of imagining does very little to render plausible the claim that Norman's belief in (2) is justified.

On the contrary, since Reliabilism predicts that Norman *is* justified in believing (2), whereas he intuitively is not, all that Williamson's example accomplishes is to add to the already rich trove of *other* Normans who serve as powerful counterexamples to Reliabilist views of justification.<sup>19</sup>

In other words, Williamson's Norman teaches us more of an old lesson about the tenability of Reliabilist views of justification than a new lesson about the shallowness of the a priori–a posteriori distinction.

I will now proceed to arguing in detail that Norman could not have come to know (2) in the way that Williamson's argument claims.

### *Using the imagination to justify belief in (2)*

Suppose Norman (who has seen nothing but red volumes of *Who's Who*) sets himself the task of imagining an instance of such a volume and comes up with an image of one that is white and decorated with a smattering of small green and yellow polka dots.

There is a perfectly good sense in which his imagination did not misfire and his performance cannot be called unskilled. In this sense, Norman is free to imagine whatever he can. We may call this sense *fanciful* imagining.

We often want imagining to be fanciful, especially when we are exploring not how things *are* but how they *might possibly have been* in the most unrestricted sense of term—metaphysically possible, as we have come to say.

How can it be that one case of imagining a false scenario is faultless, while another case of imagining a false scenario is said to be mistaken precisely because it is false? The obvious answer is that there are two different types of imaginative task here, subject to different norms. We may call this second type of imagining, *realistic* imagining as opposed to *fanciful* imagining.

Fanciful imagining is not answerable to what reality is actually like, and may not even be answerable to what reality is possibly like: Escher's drawings of geometrically impossible configurations are not faulty. Realistic imagining, on the other hand, is answerable to what reality is actually like, or to what it would be like, were certain non-actual conditions to obtain (for example, when the imagination is used to assess counterfactuals).

The kind of imagining that Norman has to engage in has to be, of course, a species of *realistic* imagining, since he is trying to use it to figure something out about the actual world.

How would Norman get himself to perform the one sort of imagining as opposed to the other? Presumably, the answer is to be found in his *intentions*. If he imagines with the intention of figuring out what is merely possible, or without any purpose at all, he imagines fancifully. If he imagines with the intention of figuring out how things are actually, or how they would have been if certain conditions had obtained, he imagines realistically.

What would it be to imagine something with the intention of using that imagining as a way of finding out about the actual world? Presumably, it would be to imagine in such a way that what one allows oneself to imagine is constrained by what one knows about the actual world. One brings to bear what one already knows, or justifiably believes, about the actual world to rule out certain imagined scenarios. When Williamson says, 'The relevant cognitive skills must be taken to include sensitivity to such matters', he is assuming that Norman is engaging in realistic imagining and constraining his imagination by his perceptually acquired background knowledge of what volumes of *Who's Who* are like.

The question is how this process is best described, and, in particular, whether it can be correctly characterized by the sort of resolutely externalist, Reliabilist picture sketched by Williamson.

In answering this, we must bear in mind that Norman is not merely using his imagination to figure out what a *particular event* was like—for example, what it was like to be at Barack Obama's first inaugural; rather, he is trying to use his imagination to figure out what property *all* the members of a given set have, on the basis of imagining a single member of that set. That is a significantly more ambitious cognitive task and carries specific epistemic burdens.

To get a better sense of what these burdens are, let's switch examples. Suppose Norman sets out to realistically imagine a VW Beetle and he imagines one accurately enough, as it happens a red one. We can't accuse Norman of imagining inaccurately, for there have been red Beetles.

But he certainly wouldn't be justified, on this basis alone, in concluding that all Beetles are red. And Norman himself, being a competent thinker, would know that. He wouldn't be tempted to conclude anything about the color of all Beetles because (we may suppose) it would be part of his background knowledge that there is no dominant color that all Beetles have.

To reasonably conclude something about the color of all Beetles from a single imagined instance, Norman would at least have to *fail* to believe that they have no dominant color. But is that enough?

Suppose Norman has no views whatsoever about whether all Beetles have the same color. He imagines a red Beetle and concludes that they all are red. I think it is clear that he would not be justified, even if in fact they were all red.

So, it looks as though, for Norman to be justified in using his single act of imagining to conclude something about the color of all the members of the set to which the imagined item belongs, he must antecedently be justified in believing that there is a single color possessed by all the members of the set, the only question being which one.

But is even this background belief sufficient to yield reasonable belief? Suppose Norman recalls a trip he took to London some years ago. He remembers being struck by the taxis there and by their distinctive dark color. But as he casts his mind back, he finds he can't recall what that color was. He remembers that all the taxis he saw had the same dark color, but for the life of him, he can't recall what that color was. Were they navy blue, or brown, or black?

I don't see Norman settling this question by thinking to himself, Let me go ahead and imagine one; surely, they are all likely to be of the color that I imagine this one to have. If Norman can't recall what color the taxis have, I don't see how his imagining one at will is going to be a good basis on which he can settle this factual question. I am not aware of any result that says that, if a subject can't recall what color F's mostly have, they are most likely to have the color that he imagines them to have, even if he's trying to imagine realistically.

Well, under what conditions *would* Norman be able to use his imagining of a single taxi to generate a justified belief about the color that all London taxis have?

It is entirely unclear to me that there are *any* conditions under which Norman would be able to do this. At a minimum, it seems to me, he must have reason to believe (a) that the taxis have a dominant color, and (b) that the taxi he imagines is appropriately representative of that dominant color.

But how could Norman possibly satisfy these two conditions without being able to *recall* what the dominant color of London taxis is? How could he know that there *is* a single color that London taxis tend to have and that *this* color is appropriately representative of what that color is, without actually being able to recall that color?

Williamson seems to think that it is enough to scotch the suggestion that Norman is *recalling*, as opposed to imagining, to stipulate that no feats of recalling particular experiences or events is involved in his description of Norman's way of coming to know (2).

But not all acts of recall need involve the recall of a particular experience or event. To advert to our earlier discussion of preservative memory, just as we may recall propositions without remembering specific events in which we learned those propositions, so we may recall generic facts about how an F looks, without recalling the look of any particular encountered F.

Having seen many London taxis, Norman may be able to recall what a London taxi looks like, without recalling any particular London taxi. He calls up a stored generic image. Having seen many frogs, Norman may be able to recall what a frog looks like, without recalling any particular frog that he has encountered. He calls up a stored generic image.

And the point is that, unless that is what is in fact going on in Norman's 'imagining' of a recent volume of *Who's Who*—unless what he's actually doing is recalling a stored generic image of a recent volume of *Who's Who*, under roughly that description—he cannot have arrived at a justified belief about all recent volumes of *Who's Who*.

I conclude that Williamson's argument's reliance on (B) leaves it open to serious doubt: Norman could not have to come to know (2) in the way that it claims.

*Using the imagination to justify belief in (1)*

What, though, about Williamson's story about how Norman comes to know (1), his claim (A)?

Williamson's account here is not easy to wrap one's mind around, because it is hard to see how someone might have learned 'crimson' and 'red' completely independently of one another.

I certainly agree that one can acquire the concepts of the primary colors independently of one another; indeed, I believe it's an important point in the theory of the a priori (to be discussed in a later piece in this volume) that one can acquire the concept *red* independently of acquiring the concept *green*. It's harder for me to see how one could acquire *red* and *crimson* without realizing that crimson is a shade of red. Borderline cases aside, having the concept *red*, one might think, is in part to have a recognitional capacity to recognize shades of red on the basis of their looks.

But let's put this worry to one side. Let's play along with the idea that Norman comes to know (1) by imagining something colored a particular shade of crimson and noting, within the imaginative supposition, that it is red, and concluding on that basis, that all crimson things are red.

Well, we saw in our discussion of (2) that if Norman is to be able to justifiably generalize on the basis of isolated acts of imagining, that he needs to satisfy certain background conditions. He needs already to know (or reasonably believe):

- (2a) That the volumes of *Who's Who* have a typical color,
- and
- (2b) That the imagined instance is appropriately representative of that typical color.

Similarly, in the case of (1) he will need to know (or reasonably believe):

- (1a) that there is a dominant color determinable to which all shades of crimson belong,
- and
- (1b) that the shade that he has imagined is properly representative in that respect of all shades of crimson.

Now, by bottom-up standards these two sets of beliefs are very different from one another: the first set could only have been justified a posteriori, while the second set would typically have been justified a priori.

Hence, under the only conditions under which Norman's beliefs could have resulted in knowledge, we see that there is a glaring difference between the two

cases that Williamson insists are epistemologically identical. Once the cases are redescribed, so that they do in fact yield knowledge on Norman's part of the relevant propositions, (D) comes out false.

Perhaps Williamson will respond that he can concede that knowledge of (1) involves a priori background beliefs, while knowledge of (2) involves a posteriori background beliefs, while insisting that the differences between them are bound to be epistemologically shallow.

In the present context, however, that answer would be question-begging. We are owed an argument that the distinction is shallow. And, so far, the only argument we have been given for that rests on claims (A), (B), (C), and (D), all of which, I have argued, are false.<sup>20</sup>

#### 4

#### *The epistemic-enabling distinction redux*

At the beginning of this paper, I rejected one attempt by Williamson to raise a theoretical difficulty for satisfactory top-down formulations of the a priori—a posteriori distinction. That difficulty turned on the claim that it is impossible to specify the notion of 'experience', in the phrase 'epistemically independent of experience', in a way that coheres with the classifications that we are prepared to make bottom-up.

My reply stated, in effect, that we could specify the notion well enough with a *list*: a justification for a belief is epistemically independent of experience provided that it does not epistemically depend on sensory experience, introspection, or substantive uses of memory. Intuition, certain uses of the imagination, and reflection could, consistently with the justification's being a priori, play an epistemic role.

Williamson has a further difficulty for putative top-down formulations that needs to be addressed, even if my reply to his challenge about the characterization of 'experience' were to be accepted.

This further difficulty, according to Williamson, is that, however exactly 'experience' is characterized, in most of its uses it will be neither purely epistemic nor purely enabling, but will somehow fall between those two poles.

Indeed, Williamson believes that that is precisely what his Norman examples illustrate. When Norman comes to 'know' that all recent volumes of *Who's Who* are red, through the use of his imagination, sense experience is not playing an evidential role—he's said not to be relying on actually perceiving such volumes or on recalling seeing such volumes. On the other hand, it's also not the case that sense experience is playing a merely enabling role, since it is doing much more than merely enabling Norman to 'access' the relevant thoughts.

Since we have rejected the claim that Norman could come to know (1) and (2) in the way claimed, we don't have much reason to worry about this particular argument.

We should point out, though, that the issue engaged here is another good example of how our topic interacts with other central issues in epistemology. For Williamson is surely right that if we were Reliabilists about justification, if we were to think of the *reliability* of our cognitive mechanisms as epistemically sufficient for the justification of the beliefs formed using them, then we would have to maintain

that sense experience may often play a role intermediate between enabling and evidential. For it will often play a role in honing the reliability of our cognitive abilities, thereby going beyond a merely enabling role, but yet fall short of playing an evidential role.

My rejection of Williamson's Norman examples is of a piece with my rejection of Reliabilist conceptions of justification: their failings, in my view, stem from the same source.

Having said that, we can hardly fault Williamson for following through on his foundational ideological commitments in epistemology and for ingeniously tracing out their consequences for the *a priori*–*a posteriori* distinction.

As a good rhetorician, Williamson tries to make us believe that the reasons he provides for his novel style of skepticism about the *a priori*–*a posteriori* distinction are independent of the ideology of Reliabilism, or the conception of justifiers as propositions. But if my counterarguments are correct, he has not succeeded in providing such ideology-independent reasons.

## 5

### *Conclusion*

What have we learned?

Our non-collusive agreement about a wide range of examples shows that the *a priori*–*a posteriori* distinction is a genuine one. And we have yet to be given a good reason for doubting that this distinction has the importance for philosophy that it appears to have.

Of course, the distinction between *a priori* and *a posteriori*, like all complicated notions in philosophy, interacts with a number of other issues—most notably, those concerning the nature of experience, independence, justification, and understanding. If you adopt certain (in my view, radical and implausible) views about these related issues, you may well end up not being able to make sense of the distinction or to find much importance in it.

But this may well show us more about your take on the related issues than about the *a priori*.

I do not want to minimize the challenges that remain for the distinction between the *a priori* and the *a posteriori*. For the moment, however, I want to emphasize that we have yet to be given a good reason for not giving the distinction the centrality in philosophy that it seems to deserve.<sup>21</sup>

## Notes

1. This paper originates as a reply to Timothy Williamson's Chapter 8, an early version of which was given at the NYIP's conference on the *a priori* held at NYU's Villa La Pietra in June of 2013. It has not previously been published.
2. I will focus on justification rather than knowledge, since I continue to regard justification as the more fundamental notion. Even in the context of this debate with Williamson, this is non-prejudicial, since Williamson is explicit that his style of skepticism about the *a priori* is not supposed to depend on his well-known 'knowledge first' conception of epistemology (see Williamson 2000).

3. For some influential criticisms of the notion of a priori justification see Quine 1951; Kitcher 1983, 2000.
4. I use the same examples that Williamson uses to illustrate a 'bottom-up' approach to the a priori–a posteriori distinction.
5. Williamson tends to say 'evidential' rather than 'epistemic'. Since evidentialism, the view that experience can only justify by providing evidence, is controversial, it is more neutral to frame the issue in the way that I have.
6. Williamson 2013: 293–4.
7. For a useful discussion of various attempts, see Casullo (2003).
8. Gibbard, private correspondence.
9. For obvious reasons, it's especially popular for those working within Bayesian frameworks.
10. See, for example, Chisholm 1977.
11. Burge 1993: 462–3.
12. Burge 1993: 465. A few paragraphs earlier, Burge allows that even substantive memories might sometimes play a non-epistemic role:
 

Substantive memories of specific events, objects, experiences . . . may play a role in deductive reasoning. They may aid reasoning without being elements in the justification they aid. So, for example, we may draw pictures in a proof, or make use of mnemonic devices to aid understanding and facilitate reasoning, without relying on them to enhance the mathematical justification. (Burge 1993: 464)

On this view, even *substantive* memories of past events, e.g., what one drew on a diagram, can play a merely enabling role, rather than an epistemic one. Sometimes Burge talks about experience 'triggering' something that makes a warrant available. I think this view is probably right, but I won't dwell on it now.
13. See Burge 1993: 462–3.
14. The thesis concerning justificational status urged in the preceding paragraphs is also supported by reflection on the well-known case of 'forgotten evidence', though I won't be able to pause to discuss it here.
15. This point is also emphasized in Casullo (2012), although we don't necessarily have the same interactions in mind.
16. See Lehrer 1990: 163–4.
17. See ch. 3 of Bonjour (1985).
18. Notice that Williamson's claim that reliability is *sufficient* for justified belief goes well beyond a claim to which some may be sympathetic, namely, that reliability is necessary for justification. I also believe, but won't argue for it here, that cases of surreptitious envatment show the necessity claim is also false.
19. See Bonjour (1985). For further critical discussion of Reliabilism see Cohen (1984) and Foley (1985).
20. Although I did not discuss (C) explicitly, its falsity follows directly from the falsity of (A) and (B).
21. This paper was initially drafted in 2012 for a session on Timothy Williamson's paper at a New York Institute of Philosophy meeting slated for November of that year. That meeting had to be postponed because of Hurricane Sandy. It was subsequently rescheduled for the summer of 2013 at NYU's Villa La Pietra, in Florence. I benefited from the comments of the audience there, and from those at the University of North Carolina at Chapel Hill, the University of Texas at Austin, and the Philosophical Society at Oxford University. I have also benefited from comments by Crispin Wright, Paul Horwich, Susanna Siegel, Antonella Mallozzi, and Jim Pryor.