

A Defense of the Significance of the A Priori – A Posteriori Distinction

Albert Casullo

University of Nebraska-Lincoln

The distinction between a priori and a posteriori knowledge has come under fire by a number of prominent authors.¹ The most sustained attack is due to Timothy Williamson. Williamson does not deny that there is a distinction. He maintains that there is a distinction but that it is superficial. Williamson initiated his attack in *The Philosophy of Philosophy* (2007) as a consequence of his novel account of knowledge of necessary truths. His argument can be summarized as follows:

- (1.1) Knowledge of necessary truths is reducible to knowledge of counterfactuals.
- (1.2) Experience can play a role in knowledge of counterfactuals that is neither purely enabling nor strictly evidential.
- (1.3) In such cases, the resulting knowledge is not happily classified either as a priori or a posteriori.
- (1.4) Therefore, the distinction between a priori and a posteriori knowledge is not useful for deep theoretical analysis.²

In a later paper, Williamson (2013) focuses his attention exclusively on the distinction.³ Here he distinguishes two approaches to introducing the distinction, bottom-up and top-down, and offers a novel argument in support of the conclusion that the distinction introduced by the bottom-up approach is superficial.

My goal is to defend four theses. First, Williamson's argument against the bottom-up distinction faces two serious objections. Its two key premises are unsupported and, moreover, the revisions necessary to support them reveal a crucial gap in the argument. Second, the argument has the unfortunate consequence of obscuring the fact that the primary issues that need to be addressed in order to assess the significance of the distinction are empirical. Third, there are compelling reasons to reject the bottom-up distinction and, moreover, Williamson's argument against that distinction cannot be extended to the top-down distinction. Fourth, although Williamson also maintains that the top-down distinction is not significant, his supporting argument is not new. It is a variation of the argument presented in Williamson (2007) and both fail for a similar reason.

1.

Williamson maintains that there are two ways of introducing the a priori – a posteriori distinction. The bottom-up way introduces the distinction by examples that are projectible. The top-down way introduces the distinction “by a direct statement of the difference between the a priori and the a posteriori in epistemologically significant theoretical terms.” (291) He points out that there are risks with each approach. The first risks introducing a distinction of little significance: “On that scenario, our classifications follow similarities and differences that, although genuine, are largely superficial, like a taxonomy of plants and animals based only on colour.” (291) Such a distinction should be avoided when theorizing because it distracts from deeper similarities and differences. The second runs the risk that all knowledge may fall on one side of the theoretical divide. In particular, it runs the risk that the concept of the a priori is vacuous.⁴ Williamson acknowledges that sophisticated attempts have been made to articulate the

distinction. Rather than address them individually, however, he proposes to address the distinction more directly.

Williamson (294) begins by outlining a general strategy for addressing the distinction more directly:

I will address the distinction more directly, by comparing what would usually be regarded as a clear case of a priori knowledge with what would usually be regarded as a clear case of a posteriori knowledge. I will argue that the epistemological differences between the two cases are more superficial than they first appear. The conclusion is not that the cases are borderline. I do not deny that they really are clear cases of a priori and a posteriori knowledge respectively, at least by bottom-up standards. . . . Rather, the appropriate conclusion is that the a priori – a posteriori distinction does not cut at the epistemological joints.

The strategy involves two steps: (1) identify a clear case of a priori knowledge and a clear case of a posteriori knowledge, and (2) argue that the epistemological differences between them are superficial.

The first step is straightforward. In order to identify clear cases of a priori knowledge and a posteriori knowledge, Williamson (295) considers two truths,

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

and argues, with respect to them, that

On the standard view, normal cases of knowledge of (1) are clearly a priori, because by definition crimson is just a specific type of red, whereas normal cases of knowledge of (2) are clearly a posteriori, because it takes direct or indirect experience of recent volumes of the British work of reference *Who's Who* to determine their colour (that is, the predominant colour of their official cover).

Williamson's starting point consists of two claims:

- (K1) Normal cases of knowledge of (1) are clearly a priori; and
- (K2) Normal cases of knowledge of (2) are clearly a posteriori.

Why accept (K1) and (K2)? Since Williamson (294) maintains that the cases in question are clear cases of a priori and a posteriori knowledge respectively by bottom-up standards, he has available a straightforward defense of (K1) and (K2) by such standards. He can argue that normal cases of knowing (1) (or relevantly similar cases) and normal cases of knowing (2) (or relevantly similar cases) are among the examples used to introduce the a priori – a posteriori distinction.⁵

The second step is more complex. In order to show that the epistemological differences between the two cases identified in step 1 are superficial, Williamson offers a more detailed description of each case. The description of the first case involves a subject, Norman, an

explanation of how he learns the expressions ‘crimson’ and ‘red’ and, finally, an account of how he comes to know (1). The account of how Norman comes to know (1) consists of two steps followed by an observation regarding the cognitive capacity underlying the second step:

- (A1) First, Norman uses his skill in making visual judgments with ‘crimson’ to visually imagine a sample of crimson.
- (A2) Then he uses his skill in making visual judgments with ‘red’ to judge, within the imaginative supposition, ‘It is red’.
- (A3) This involves a general human capacity to transpose ‘online’ cognitive skills originally developed in perception into corresponding ‘offline’ cognitive skills subsequently applied in imagination. (295)

Since Norman’s performance was sufficiently skillful, Williamson maintains that he knows (1).

The description of the second case involves the same subject, Norman, an explanation of how he learns the expressions ‘recent volumes of *Who’s Who*’ and ‘red’ and an account of how he comes to know (2). The account of how Norman comes to know (2) also consists of two steps followed by an observation regarding the cognitive capacity underlying the second step:

- (B1) 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*.
- (B2) Then he uses his skill in making visual judgments with ‘red’ to judge, within the imaginative supposition, ‘It is red’.

- (B3) 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. (296)

Once again, Williamson maintains that since Norman’s performance was sufficiently skillful, he knows (2).

On the basis of the more detailed descriptions of the two cases, Williamson (296-297) offers the following argument against the distinction:

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, the a priori – a posteriori distinction is epistemologically shallow.

Williamson’s contention that there are no significant epistemological differences between Norman’s knowledge of (1) and Norman’s knowledge of (2) rests on the psychological claim that the cognitive processes underlying both are almost exactly similar. The psychological claim is supported by his description of the two cases. In particular (A3), which is his description of the cognitive capacity underlying Norman’s knowledge of (1), is identical to (B3), which is his description of the cognitive capacity underlying Norman’s knowledge of (2).

Williamson’s argument can be articulated as follows:

- (2.1) Norman's knowledge of (1) is clearly a priori.
- (2.2) Norman's knowledge of (2) is clearly a posteriori.
- (2.3) The cognitive processes underlying Norman's knowledge of (1) are almost exactly similar to those underlying Norman's knowledge of (2).
- (2.4) Therefore, the a priori – a posteriori distinction is shallow.

Since Williamson regards his descriptions of the two Norman cases to be more detailed descriptions of a normal case of a priori knowledge and a normal case of a posteriori knowledge, he takes premises (2.1) and (2.2) to be straightforward consequences of (K1) and (K2) respectively. His detailed descriptions of the cognitive processes underlying Norman's knowledge of (1) and Norman's knowledge of (2) support premise (2.3).

The argument is flawed. It is open to two objections. Premises (2.1) and (2.2) are unsupported.⁶ Moreover, the revisions necessary to support them reveal a significant gap in the argument. Consider premises (2.1) and (2.2). Despite Williamson's assumption to the contrary, (2.1) does not follow from (K1) and (2.2) does not follow from (K2). The source of the problem is a critical omission in Williamson's starting point. He does not specify what constitutes a "normal case" of knowledge of (1) or what constitutes a "normal case" of knowledge of (2). There are two issues. First, Williamson does not offer a general characterization of what he means by a "normal case" of knowing p. Presumably, a normal case of knowing p is some particular way of knowing p. But what are the *general features* that distinguish the normal way of knowing p from other ways of knowing p? Is the normal way of knowing p the predominant or typical way of knowing p in some epistemic community? Is it the way in which the experts in

an epistemic community know p ? Or is it some other way? Second, he does not indicate *which particular way* of knowing (1) is the normal way of knowing (1) or *which particular way* of knowing (2) is the normal way of knowing (2). Although he observes that by definition crimson is just a type of red, this observation does not indicate which particular way of knowing (1) constitutes a normal case of knowing (1). Similarly, although he observes that it takes direct or indirect experience of recent volumes of *Who's Who* to determine their color, this observation does not indicate which particular way of knowing (2) constitutes a normal case of knowing (2).

This omission bears directly on (2.1) and (2.2). Norman's way of knowing (1), according to Williamson, is via an imaginative exercise constituted by steps (A1) - (A2). (2.1) follows from (K1) only if Norman's way of knowing (1) is the normal way of knowing (1). But Williamson has offered no basis for the claim that Norman's way of knowing (1) is the normal way of knowing (1) since he has neither indicated which particular way of knowing (1) is the normal way of knowing (1) nor provided any general criteria for determining whether Norman's way of knowing (1) is the normal way. He simply assumes that (A1) - (A2) constitute the normal way of knowing (1). Similarly, (2.2) follows from (K2) only if Norman's way of knowing (2) is the normal way of knowing (2). But Williamson has offered no basis for the claim that Norman's way of knowing (2) is the normal way of knowing (2). He simply assumes that (B1) - (B2) constitute the normal way of knowing (2).⁷

My initial objection to Williamson's argument may appear to be inconsequential since there are two obvious ways to revise the argument in response to it. The first is to drop all talk of "normal cases" of knowledge and to replace (K1) and (K2) respectively with:

(K1*) *All* ways of knowing (1) are clearly a priori; and

(K2*) *All* ways of knowing (2) are clearly a posteriori.

Since (2.1) and (2.2) are straightforward consequences of (K1*) and (K2*) respectively, my concern about what constitutes a normal case of knowledge of (1) and what constitutes a normal case of knowledge of (2) is irrelevant to the revised argument. The second way of revising the argument is to add to the original argument the following premises:

(K1+) There is *only one* way of knowing (1).

(K2+) There is *only one* way of knowing (2).

If there is only one way of knowing (1) and only one way of knowing (2), my concern about which particular way of knowing (1) and which particular way of knowing (2) is the normal way is irrelevant to the revised argument.

The proposed revisions reveal a significant gap in Williamson's argument. He fails to take into account the possibility of epistemic overdetermination.⁸ There are two types of epistemic overdetermination that bear on his argument, one general and one more specific:

(GO) One can know *p* in more than one way.

(SO) One can know *p* in more than one way, some a priori and some a posteriori.

(K1+) is incompatible with (GO); (K1*) is incompatible with (SO). Moreover, his failure to do so is puzzling since he (292) endorses both (GO) and (SO):

The distinction between the a priori and the a posteriori is primarily a classification of specific ways of knowing. . . . One may know *p* both a priori and a posteriori, if one knows it in several ways, some a priori, some a posteriori. Tradition excluded that case But that was a mistake.

Since Williamson acknowledges the general possibility that one can know p both a priori and a posteriori, in the absence of some further argument, he has no basis for denying that one can know (1) both a priori and a posteriori.

How does this point bear on his argument? Consider again Williamson's account of how Norman comes to know (1) and his account of how Norman comes to know (2). Suppose that we grant that Williamson is right about the cognitive processes underlying both Norman's knowledge of (1) and Norman's knowledge of (2). It does not follow that there is no deep epistemological distinction between knowing (1) and knowing (2). Why not? Suppose that there is an alternative way of knowing (1), such as via rational insight. Since Williamson has not shown that the cognitive processes underlying the alternative way of knowing (1) and those underlying knowing (2) are almost exactly similar, it remains open that they differ in some epistemologically significant way.⁹ But if they do differ in some epistemologically significant way, then it is false that the a priori – a posteriori distinction is shallow.

The gap in Williamson's argument has significant repercussions for his case against the significance of the a priori – a posteriori distinction. Given that no contemporary proponent of the a priori embraces Williamson's account of the cognitive processes underlying such knowledge, they all can accept his account of Norman's knowledge of (1) and Norman's knowledge of (2), yet deny his conclusion that the a priori – a posteriori distinction is epistemologically shallow. Moreover, they are in a position to turn his argument against him. They can offer the following counterargument:

- (3.1) Knowledge of (1) is overdetermined: there are (at least) two ways of knowing (1), (at least) one is a priori and (at least) one is a posteriori.
- (3.2) The a posteriori way of knowing (1) employs cognitive processes quite *similar* to those employed by Norman's way of knowing (2).
- (3.3) The a priori way of knowing (1) employs cognitive processes quite *different* from those employed by Norman's way of knowing (2).
- (3.4) Therefore, there is a deep epistemological distinction between the a priori and the a posteriori.

The import of this counterargument goes beyond Williamson's treatment of knowledge of (1). The broader aim of his paper is to argue that knowledge of truths of set theory, mathematics and logic involves the same cognitive processes as Norman's knowledge of (1). But, once again, this establishes at most that there is *a* way of knowing these truths that employs cognitive processes similar to those involved in Norman's knowledge of (2). It does not establish that there is *no* way of knowing those truths that employs different cognitive processes. Hence, the proponent of the a priori can accept everything that he says about knowledge of truths of set theory, mathematics and logic, but deny his conclusion that the a priori – a posteriori distinction is shallow.

In order to sustain his conclusion, Williamson must deny that there are alternative a priori ways of knowing (1) and the truths of set theory, mathematics and logic. He cannot do so without engaging the claims of proponents of the a priori to the contrary. A denial that does not engage those claims begs the question. Consequently, Williamson cannot establish that the a

priori – a posteriori distinction is insignificant without engaging substantive issues regarding the source of such knowledge. Stipulations about the cognitive processes underlying the normal way of knowing (1) and the truths of set theory, mathematics and logic are insufficient.

2.

Although the argument fails and, as a result, is not a threat to the a priori – a posteriori distinction, it has an unfortunate consequence. It obscures the primary issues that need to be addressed in order to assess the significance of the a priori – a posteriori distinction. As the lacuna in Williamson's argument clearly reveals, the primary issues that need to be addressed are *empirical*. Williamson's argument rests on two pairs of empirical claims:

(E1a) Norman's belief that (1) is based on an exercise in imagination that employs offline cognitive skills developed online in perception.

(E1b) Norman's exercise in imagination is sufficiently reliable that the resulting belief constitutes knowledge.

(E2a) Norman's belief that (2) is based on an exercise in imagination that employs offline cognitive skills developed online in perception.

(E2b) Norman's exercise in imagination is sufficiently reliable that the resulting belief constitutes knowledge.

It also involves a tacit empirical assumption:

(EA) There is no way of knowing (1) that employs cognitive processes quite *different* from those employed by Norman's way of knowing (2).

Why accept these empirical claims? In the case of the (E1a) and (E2a) pair, Williamson (296) acknowledges that he has offered only a "broad-brush description" that "neglects many issues." Nevertheless, he (296) also contends that "we may accept it as a sketch of the cognitive processes" underlying both Norman's knowledge of (1) and Norman's knowledge of (2). Williamson, however, offers no evidence in support of that claim. It appears to be based entirely on his reflections from the armchair regarding how he acquires such beliefs. Such armchair reflection, however, does not constitute compelling evidence for a claim in cognitive psychology. At best, such reflection might yield insightful hypotheses about the underlying cognitive processes involved in coming to accept propositions such as (1) and (2). The hypotheses, however, cannot be confirmed on the basis of anecdotal evidence from the armchair. Empirical investigation is necessary.

In fairness to Williamson, it should be acknowledged that proponents of the a priori also employ his method. The key premise of their response to his argument is

(3.3) The a priori way of knowing (1) employs cognitive processes quite *different* from those employed by Norman's way of knowing (2).

But how do they support this contention? Typically they reflect on the way in which they come to believe propositions such as (1) and claim that such reflection reveals that the cognitive

process in question involves intuition or rational insight, which are alleged to be distinctive cognitive processes that are not involved in coming to believe propositions such as (2).¹⁰ But, once again, armchair reflection is not sufficient to establish conclusions about the underlying cognitive processes involved in belief formation.

(E1b) and (E2b) introduce empirical claims of a different sort. Assume that Williamson has correctly identified the cognitive processes employed in normal cases of believing (1) and (2). Why think that they are reliable? Williamson (297) offers two responses. First, he maintains that to deny that they are reliable “imposes an idealized epistemological standard for knowledge that humans cannot be expected to meet.” Second, he contends that “even if we are skeptical about knowledge in such cases, we should still assign belief in (1) and (2) some other sort of positive epistemic status, such as reasonableness, to which the a priori – a posteriori distinction should still apply in some form.” Both responses are off the mark. Both attempt to settle an empirical question about the reliability of a belief forming process on the basis of broadly a priori philosophical considerations.

With respect to the first response, if results in cognitive science establish that imaginative exercises of the type that Williamson envisages are quite unreliable, it does not follow that cognitive science imposes some idealized standard for human knowledge. Whatever standards for human knowledge that Williamson articulates, insofar as they involve considerations of reliability, it remains an open empirical question whether any particular belief-forming process meets that standard. No appeal to some perceived threat of skepticism can settle the issue. With respect to the second response, even if the standards for reasonable belief require a lower degree of reliability than the standards for knowledge, it still remains an open empirical

question that cannot be settled with anecdotal armchair evidence whether any particular belief-forming process meets that lower standard.

(EA) introduces an additional empirical claim: Either Norman's way of knowing (1) is the *only* way of knowing (1) or any other way of knowing (1) employs cognitive processes quite *similar* to those employed by Norman's way of knowing (2). Armchair reflection cannot establish either that Norman's way of knowing (1) is the *only* way of knowing (1) or that any other way of knowing (1) employs cognitive processes quite *similar* to those employed by Norman's way of knowing (2). After all, proponents of the a priori contend, on the basis of such reflection, that there are ways of knowing (1) that employ cognitive processes quite different from those employed by Norman's way of knowing (2). Empirical investigation is necessary to adjudicate these competing claims.

Williamson (309) acknowledges that work in cognitive psychology is relevant to epistemology: "Cognitive psychology will have much to offer the epistemologist's attempt to overcome philosophical prejudices and classify according to deeper and less obvious similarities and differences." Cognitive psychology, however, has much more to offer the epistemologist than he acknowledges. The results of cognitive psychology are necessary to address each of the following issues:

- (II) Are normal cases of believing (1) and (2) based on exercises in imagination that employ offline cognitive skills developed online in perception?

- (I2) If normal cases of believing (1) and (2) are based on exercises in imagination, are such exercises sufficiently reliable to bestow positive epistemic status of any sort on the beliefs that they produce?¹¹
- (I3) If normal cases of believing (1) and (2) are based on exercises in imagination, are there ways of coming to believe (1) that employ cognitive processes quite *different* from such exercises in imagination?
- (I4) If there are alternative ways of coming to believe (1), are they sufficiently reliable to bestow positive epistemic status of any sort on the beliefs that they produce?

In sum, Williamson's discussion obscures the important empirical questions that need to be addressed in order to determine whether the a priori – a posteriori distinction is deep and underestimates the role of cognitive psychology in answering those questions.

3.

Williamson contends that there are two approaches to introducing the a priori – a posteriori distinction: top-down and bottom-up. The second plays an essential role in his argument in support of the claim that the distinction is superficial. Williamson appeals to it in support of premises of (2.1) and (2.2). There are, however, compelling reasons to reject that approach to introducing the distinction.

Williamson's discussion of the bottom-up approach is sketchy. The central idea is that when the distinction is introduced bottom-up, it is introduced by examples. The introduction involves two steps. We initially notice certain obvious similarities and differences among the examples and then project the distinction to new cases by the presence or absence of these

similarities or differences. The distinction, made in this way, runs the risk of superficiality because the obvious similarities and differences may turn out to be superficial from a theoretical perspective.

What are the similarities or differences relevant to the introduction of the a priori – a posteriori distinction? Williamson (298) offers the following suggestion:

The most salient difference between Norman's knowledge of (2) and his knowledge of (1) is that (2) is contingent while (1) is necessary. That difference may be what inspires the idea that there must be a deep difference in his knowledge of them.

But, as Williamson points out, Kripke has shown that difference in metaphysical status need not entail a difference in epistemological status. Nevertheless, Williamson (300) maintains:

Even having accepted Kripke's examples of the contingent a priori and the necessary a posteriori, we may still operate on the default assumptions that knowledge of necessary truths is a priori and that knowledge of contingent ones is a posteriori. Unlike Kripke's, cases such as Norman's may trigger nothing to overturn either default, especially when they are considered separately, so by default we confidently classify knowledge of the necessary truth as a priori and knowledge of the contingent one as a posteriori, without noticing that there is no significant *epistemological* difference between them.

Hence, the obvious similarities and differences that we notice and project when introducing the a priori – a posteriori distinction by the bottom-up approach are similarities and differences in metaphysical (or modal) status.

Williamson's observations on the bottom-up approach provide a compelling reason to reject it. The bottom-up approach does not introduce an epistemological distinction; it introduces a metaphysical distinction. Consequently, any argument based on premises supported by the bottom-up approach that draws conclusions about the import of the a priori – a posteriori distinction should be rejected. Such an argument only establishes that the necessary – contingent distinction is epistemologically shallow. It does not establish that the relevant epistemological distinction is epistemologically shallow.

There is also a more general theoretical reason for rejecting the bottom-up approach to introducing the distinction. There are two central questions regarding a priori knowledge. What is a priori knowledge? Is there a priori knowledge? The primary theoretical motivation for addressing the first is to answer the second. The bottom-up approach to addressing the first question, however, appears to trivially entail an affirmative answer to the second.¹² If the bottom-up approach provides an adequate response to the first question, then it follows that there are clear cases of a priori knowledge. Consequently, the distinction introduced by the bottom-up approach cannot nontrivially address the question that motivates the introduction of the distinction.

Williamson maintains that there is a clear a priori – a posteriori distinction by bottom-up standards. That distinction, he argues, is superficial because it masks important epistemological similarities. The distinction introduced by bottom-up standards masks important epistemological

similarities because it tracks differences in metaphysical status rather than differences in epistemological status. But there are good reasons to reject the distinction introduced by bottom-up standards. Moreover, Williamson's argument in support of the claim that the a priori – a posteriori distinction is superficial depends essentially on the bottom-up approach to introducing the distinction. It cannot be extended to the distinction introduced by top-down standards.

In order to extend that argument to show that the top-down distinction is superficial, Williamson would have to show that Norman's knowledge of (1) is clearly a priori and that Norman's knowledge of (2) is clearly a posteriori by top-down standards. The prospects of doing so are dim since the conjunction of (2.1*), (2.2*) and (2.3) is an inconsistent triad:

(2.1*) Norman's knowledge of (1) is clearly a priori by top-down standards.

(2.2*) Norman's knowledge of (2) is clearly a posteriori by top-down standards.

(2.3) The cognitive processes underlying Norman's knowledge of (1) are almost exactly similar to those underlying Norman's knowledge of (2).

Either the cognitive processes underlying Norman's knowledge of (1) and Norman's knowledge of (2) are (epistemically) dependent on experience or they are not. If they are, then both Norman's knowledge of (1) and Norman's knowledge of (2) are a posteriori. If they are not, then both Norman's knowledge of (1) and Norman's knowledge of (2) are a priori.

In short, Williamson's argument in support of the claim that the a priori - a posteriori distinction is superficial is of limited import. It depends essentially on the bottom-up approach, which introduces a distinction that tracks differences in metaphysical status rather than

differences in epistemological status. His argument, however, leaves untouched the top-down approach, which introduces a distinction that tracks differences in epistemological status.

4.

Williamson's leading argument is directed at the distinction introduced by the bottom-up approach. The distinction introduced by the top-down approach, however, does not escape scrutiny. Williamson maintains that it faces a serious problem.¹³ It must articulate the theoretical epistemological terms that it employs to draw the a priori – a posteriori distinction. According to Williamson (292), "A way of knowing is a priori if and only if it is independent of experience." The top-down approach must articulate the relevant sense of the terms 'independent' and 'experience'. With respect to the former, Williamson (293) maintains that the standard response is to maintain that "an a priori way of knowing may depend on experience in its enabling role but must not depend on experience in its evidential role." A problem, however, now emerges for the top-down approach:

The top-down way of introducing the distinction between the a priori and the a posteriori promised to put it on a firm theoretical footing, but in practice relies on other terms (such as 'experience' and 'evidence') understood at least partly bottom-up, through examples and prototypes. Although bottom-up understanding often serves us well enough, in the present case it leaves us puzzled all too soon. (294)

Williamson (298) contends that an examination of the Norman cases supports this contention since it reveals that "the role of experience in both cases is more than purely enabling but less

than strictly evidential,” which “reinforces a suspicion raised in section 2, that talk of ‘experience’ and ‘evidence’ does little to help us apply the a priori – a posteriori distinction top-down.”

The argument against the top-down approach presented by the two Norman cases can be summarized as follows:

- (4.1) According to the top-down approach, a way of knowing is a priori if and only if it does not depend on experience in its evidential (as opposed to enabling) role.
- (4.2) Experience plays a role in both Norman’s knowledge of (1) and Norman’s knowledge of (2) that is more than purely enabling but less than strictly evidential.
- (4.3) Our bottom-up understanding of the terms ‘experience’ and ‘evidence’ does not help us apply the a priori – a posteriori distinction top-down in these cases.
- (4.4) Therefore, we are left puzzled as to whether Norman’s knowledge of (1) and Norman’s knowledge of (2) is a priori or a posteriori.
- (4.5) Therefore, the a priori – a posteriori distinction (introduced top-down) is not theoretically significant.

The top-down approach to introducing the distinction requires new theoretical epistemological terms. Unless the new theoretical terms can be articulated, the top-down introduction is not illuminating. The new terms can be articulated either top-down or bottom-up. Our understanding of the new terms is bottom-up, according to Williamson, but the Norman cases show that the bottom-up understanding is inadequate for theoretical purposes.

This argument is not new. It is a variation of the argument presented in Williamson (2007). The leading premise in both arguments is the same: there are cases of knowledge in which experience plays a role that is more than purely enabling but less than strictly evidential. The novel aspect of the present version of the argument is Williamson's contention that our bottom-up understanding of the terms 'experience' and 'evidence' leaves us puzzled as to whether such cases of knowledge are a priori or a posteriori.

Williamson's diagnosis of the source of our puzzlement in cases where experience plays a role that is more than enabling but less than evidential is mistaken. He maintains that our puzzlement is due to the fact that our understanding of the terms 'evidence' and 'experience' is bottom-up. This understanding serves us well enough in classifying familiar cases but fails us in cases of knowledge where experience plays a role that is more than purely enabling but less than evidential. Hence, according to Williamson, our puzzlement over the Norman cases arises because either (P1) or (P2) is true:

(P1) When we consider the two Norman cases, we are puzzled as to whether the term 'experience' applies.

(P2) When we consider the two Norman cases, we are puzzled as to whether the term 'evidence' applies.

(P1) and (P2), however, are false.¹⁴

According to (P1), when we consider the Norman cases, we are puzzled as to whether they are cases of a priori knowledge because the top-down approach states that a way of knowing

is a priori if and only if it is independent of experience, but it is unclear or indeterminate whether the term ‘experience’ applies to some feature of the cases on which, according to Williamson, Norman’s knowledge depends. (P1) is false since the only epistemically relevant features involved in Williamson’s description of the two Norman cases are the perceptual experiences necessary to develop Norman’s capacity to make reliable visual judgments with the expressions ‘red’, ‘crimson’ and ‘recent volumes of *Who’s Who*’. But it is not unclear or indeterminate whether the term ‘experience’ in the top-down definition of a priori knowledge applies to perceptual experiences. It is uncontroversial that dependence on perceptual experience is incompatible with a priori knowledge.¹⁵

In order to avoid misunderstanding of my contention regarding (P1), it is important to recognize that there is a significant problem regarding our understanding of the term ‘experience’ as it applies to a priori knowledge, but that problem does not arise when we consider the two Norman cases. What is the problem? It is generally accepted that an a priori way of knowing is one that is independent of experience. Proponents of the a priori, however, maintain that there is a sense of the term ‘experience’ on which an a priori way of knowing *can* depend on experience. Consider the following passage from Bonjour (1998, 7):

It is obvious at once that the broadest meaning of the term ‘experience’, that in which it refers to any sort of mental process that one consciously undergoes, is substantially too broad; in that sense, following a mathematical proof or even reflecting on a supposedly self-evident proposition would be an instance of experience, and *a priori* justification would be ruled out in a trivial and uninteresting way.

Let us call the sense of ‘experience’ that applies to any sort of conscious process (or state) that one undergoes (instantiates) the *broad* sense. The sense of ‘experience’ relevant to articulating the a priori – a posteriori distinction is narrower than the broad sense.

There is another sense of ‘experience’ that is narrower than the broad sense. This sense of the term ‘experience’ applies to sense experience in its various modes; that is, to the experiences associated with the five senses. Call this the *narrow* sense of experience. But, as Bonjour (1998, 7) contends, this is not the sense of ‘experience’ relevant to articulating the a priori – a posteriori distinction:

But it is just as obvious that the relevant concept of experience cannot be confined in its scope to the obvious paradigm of such experience, namely, the experience involved in ordinary sense-perception and deriving from the five standard senses. The justification of introspective knowledge pertaining to one’s own states of mind should surely count as empirical, as should that of kinesthetic knowledge of the position and movements of one’s body and that of knowledge of past events deriving, via memory, from previous episodes of perception.

Hence, the sense of ‘experience’ relevant to articulating the a priori – a posteriori distinction is broader than the narrow sense but narrower than the broad sense. It includes, but is not limited to, the experiences associated with the five senses. Let us call this sense the *intermediate* sense.

There are two general approaches to articulating the intermediate sense of ‘experience’. The first to offer an exhaustive list of experiential sources of knowledge (justification). The

second is to offer a general characterization of the difference between experiential and nonexperiential sources of knowledge (justification). The first is analogous to Williamson's bottom-up approach; the second is analogous to his top-down approach. I (Casullo 2003, 2013) have argued that both approaches face significant challenges. There is disagreement over which sources count as experiential sources, and all the extant proposals for a general characterization are open to objection. So there is a significant problem regarding our understanding of the intermediate sense of 'experience'. Hence, it might appear that the fact that there is a significant problem understanding the intermediate sense of 'experience' lends support to (P1).

Although there is a significant problem understanding the intermediate sense of 'experience', that problem does not lend support to (P1). The problem we face in understanding the intermediate sense of 'experience' is in understanding to which experiences it applies *in addition to* the experiences associated with the five senses. But the only epistemically relevant experiences involved in Williamson's description of the two Norman cases are perceptual experiences. And it is obvious that the intermediate sense of 'experience' applies to perceptual experiences.

(P2) is also false. In order to see why, it is important to note that Williamson maintains that the two Norman cases can be interpreted in (at least) two different ways. Consider Norman's knowledge of (1). Here Williamson (298) maintains:

One interpretation of the example is that, although Norman's knowledge of (1) does not depend on episodic memory, and he may even lack all episodic memory of any relevant particular colour experiences, he nevertheless retains from such experiences generic

factual memories of what crimson things look like and of what red things look like, on which his knowledge of (1) depends. . . . On this interpretation, Norman's colour experience plays an evidential role in his knowledge of (1), thereby making that knowledge a posteriori.

Norman's knowledge of (2) can also be interpreted in this way.

Williamson (298), however, also points out that the two Norman cases can be interpreted in a different way:

Instead, it may be proposed, although colour experience can play an evidential role in a posteriori knowledge of what crimson things look like, and so indirectly in a posteriori knowledge of (1), we need not develop the example that way. The only residue of Norman's colour experience active in his knowledge of (1) may be his skill in recognizing and imagining colours. Such a role for experience, it may be held, is less than strictly evidential.

Norman's knowledge of (2) can also be interpreted in this way. So we have two different interpretations of the Norman cases. On one interpretation, Norman's perceptual experiences play an (indirect) evidential role but, on the other, Norman's perceptual experiences do *not* play an evidential role.

One might suggest that, by acknowledging that there are these two different interpretations of the two Norman cases, Williamson is providing support for (P2). The suggestion is that, when

we consider the two Norman cases, we are puzzled as to whether the term ‘evidence’ applies because we are sensitive to the fact that the cases can be interpreted in the two different ways that Williamson acknowledges and it is unclear or indeterminate which is the relevant interpretation. This suggestion, however, misses the point of Williamson’s introduction of the two interpretations of the Norman cases. The point of introducing the two interpretations is to make clear that he intends the cases to be interpreted in the second way; that is, he intends the cases to be interpreted in such a way that experience is not playing an evidential role. Moreover, if it were unclear or indeterminate whether the term ‘evidence’ applied to the two Norman cases, then premise (4.2) of his argument would be compromised. Hence, contrary to the suggestion, Williamson’s argument against the top-down approach presupposes that (P2) is false.

Williamson’s argument against the top-down approach to articulating the a priori – a posteriori distinction rests on two claims:

- (C1) We are puzzled as to whether cases of knowledge in which experience plays a role that is more than purely enabling but less than strictly evidential should be classified as a priori or a posteriori.
- (C2) This puzzlement is due to the fact that the top-down articulation of the concept of a priori knowledge involves the terms ‘experience’ and ‘evidence’, which are introduced bottom-up by examples, and it is unclear or indeterminate whether the terms apply to such cases.

I have argued that (C2) is false: Williamson's diagnosis of the source of our puzzlement is mistaken. This does not vindicate the top-down approach unless we can offer an alternative explanation of the source of our puzzlement. I will conclude by arguing that Williamson's puzzle is generated by his own epistemological views. It results from the fact that the top-down articulation of the concept of a priori knowledge that he offers in (4.1) does not fit well with his general theory of knowledge. The mismatch is the source of our puzzlement. As we will see, once we identify the mismatch and modify his articulation of the concept of a priori knowledge to eliminate it, the puzzle vanishes.

My contention is that Williamson's puzzlement arises from offering an articulation of the concept of a priori knowledge that does not cohere well with his general theory of knowledge. In order to bring out this point more clearly, let us briefly consider the example that Williamson (2007) employs to support premise (1.2) of his earlier argument. Consider a person who learns the words 'inch' and 'centimeter' independently of one another by learning to make reliable naked eye judgments of distances. Williamson (2007, 166) contends that such a person can employ his or her capacity to judge distances visually offline to make the following counterfactual judgment:

- (25) If two marks had been nine inches apart, they would have been at least nineteen centimeters apart.

The person can do so by visually imagining two marks nine inches apart and employing his or her capacity to judge distances in centimeters visually offline to judge that they are at least nineteen centimeters apart. If the person's judgment is sufficiently reliable, that person knows (25). The

role of experience in this case is more than purely enabling but it is not evidential. Williamson (2007, 166) articulates its role as follows:

I know (25) only if my offline application of the concepts of an inch and a centimeter was sufficiently skilful. Whether I am justified in believing (25) likewise depends on how skilful I am in making such judgments. My possession of the appropriate skills depends constitutively, not just causally, on past experience for the calibration of my judgments of length in those units.

The central point for our purposes is that both knowledge that (25) and justified belief that (25) depends epistemically on the skillful application of the relevant concepts which, in turn, depends constitutively on past experience.

The traditional understanding of the concept of a priori knowledge is top-down. I (Casullo 2003) defend the following (top-down) articulation of the traditional concept:

(APK) S knows a priori that p if and only if S's belief that p is justified a priori and the other conditions on knowledge are satisfied; and

(APJ) S's belief that p is justified a priori if and only if S's justification for the belief that p does not depend on experience.

The traditional concept of a priori knowledge is embedded within a more general epistemological framework that maintains that justification is a necessary condition for knowledge. Although that

condition has been challenged in contemporary epistemology, it remains widely accepted.

Moreover, Williamson appears to endorse it.

There is also controversy in contemporary epistemology over how to articulate the concept of justification. One controversy is over the relationship between the concept of justification and the concept of evidence. Two points are important for our purposes. First, within a general theory of knowledge that endorses

(JE) Justification is a function of evidence alone,¹⁶

(APJ) reduces to

(APE) S's belief that p is justified a priori if and only if S's evidence for the belief that p does not depend on experience.

Within such a theory of knowledge, (APE) is a natural way of articulating the traditional concept of a priori knowledge.¹⁷ It is an articulation of the concept that coheres with the more general epistemological commitments of that theory. Second, within a theory of knowledge that rejects (JE), (APE) is not a natural way of articulating the traditional concept of a priori knowledge. It is an articulation that does not cohere well with the more general epistemological commitments of that theory. The general theory of knowledge in which (APE) is embedded maintains that justification is not a function of evidence *alone* but (APE) maintains that justification is a priori if it is independent of experiential evidence *alone*. Hence, (APE) fails to take into account the

nonevidential features relevant to justification and, as a consequence, gives rise to the following puzzling question:

- (PQ) If S's justification for the belief that p does not depend on any experiential evidence but does depend (epistemically) on some nonevidential experiential feature, is S's belief that p justified a priori or a posteriori?

According to (APE), S's belief is justified a priori but, according to the general theory of knowledge in which (APE) is embedded, S's belief is justified (in part) by experience. Hence, we are left puzzled.

The puzzle that Williamson identifies in (4.4) of his argument has the same source. It results from the conjunction of his endorsement of (APE) and his denial of (JE). Premise (4.1) represents an endorsement of (APE) since it entails that if S's justification for the belief that p does not depend on experiential evidence then S's justification is a priori.¹⁸ It is silent about the nonevidential features relevant to justification. Premise (4.2) entails the denial of (JE) since it entails that knowledge (and, presumably, justification) can depend epistemically on experience even if it does not depend on experience in its evidential role. Hence, Williamson's diagnosis of the source of the puzzle is mistaken. The source of the puzzle is not the bottom-up understanding of 'experience' and 'independence' or, more generally, the traditional concept of a priori knowledge. The source of the puzzle is that he offers an articulation of the concept of a priori knowledge that does not cohere well with the general features of the theory of knowledge that he endorses.

If (APE) is replaced by (APJ), Williamson's puzzle vanishes since the conjunction of (APJ) and (4.2) straightforwardly entails the conclusion:

(4.4*) Therefore, both Norman's knowledge of (1) and Norman's knowledge of (2) are a posteriori.

Our bottom-up understanding of the terms 'experience' and 'evidence' poses no obstacle to drawing this conclusion. Moreover, the conclusion should not be surprising or controversial. If Williamson is correct in maintaining that the cognitive processes underlying Norman's knowledge of (1) and Norman's knowledge of (2) are almost exactly similar and that both depend epistemically on experience, then both are known a posteriori.

5.

Williamson maintains that the a priori – a posteriori distinction is superficial. He notes that there are two ways of introducing the distinction: bottom-up and top-down. His leading argument is directed against the distinction introduced by the bottom-up approach. I contend that the leading premises of the argument, (2.1) and (2.2), are unsupported and that, even if they are granted, the conclusion of the argument does not follow. I go on to argue that there are compelling reasons to reject the bottom-up approach to introducing the distinction.

Williamson, however, maintains that the top-down approach also faces a serious problem since it involves the terms 'experience' and 'evidence', which are understood bottom-up. That approach leaves us puzzled over how to apply the distinction to cases of knowledge where experience plays a role that is more than purely enabling but less than strictly evidential. I argue

that our puzzlement about these cases is not due to our understanding of the terms ‘experience’ and ‘evidence’. It is due to the fact that Williamson offers a top-down articulation of the concept of a priori knowledge that does not cohere well with his general theory of knowledge.

Although Williamson’s challenge to the a priori – a posteriori distinction fails, there remains a deeper challenge. Since there are compelling reasons to reject the bottom-up approach to introducing the distinction, the only remaining option (short of maintaining that it is primitive) is to introduce it top-down. But, as Williamson notes, this approach introduces the terms ‘independent’ and ‘experience’, which require articulation. The articulation of ‘independent’ has received considerable attention. Some, such as myself (Casullo 2003, 2009), favor

(IN1) S’s belief that p is justified a priori if and only if S’s belief that p is nonexperientially justified.

Others, such as Kitcher (1983, 2000), favor

(IN2) S’s belief that p is justified a priori if and only if S’s belief that p is nonexperientially justified and cannot be defeated by experience.

The articulation of ‘experience’, however, has received far less attention. In particular, the challenge, introduced in section 4, of articulating the intermediate sense of ‘experience’ remains open. If this challenge cannot be met, then there are compelling grounds for questioning the coherence (and not merely the significance) of the a priori – a posteriori distinction.^{19, 20}

References

- Bealer, George. 1992. "The Incoherence of Empiricism." *Proceeding of the Aristotelian Society* supp. vol. 66: 99-138.
- BonJour, Laurence. 1998. *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- Brueckner, Anthony. 2011. "Albert Casullo's *A Priori Justification*." In *What Place for the A Priori?*, eds. M. J. Shaffer and M. L. Veber. Chicago: Open Court.
- Casullo, Albert. 2003. *A Priori Justification*. New York: Oxford University Press.
- _____. 2005. "Epistemic Overdetermination and A Priori Justification." *Philosophical Perspectives* 19: 41-58.
- _____. 2009. "Analyzing A Priori Knowledge." *Philosophical Studies* 142: 77-90.
- _____. 2011. "Reply to my Critics: Anthony Brueckner and Robin Jeshion." In *What Place for the A Priori?*, eds. M. J. Shaffer and M. L. Veber. Chicago: Open Court.
- _____. 2012a. "Articulating the A Priori - A Posteriori Distinction." In *Essays on A Priori Knowledge and Justification*. New York: Oxford University Press.
- _____. 2012b. "Counterfactuals and Modal Knowledge." In *Essays on A Priori Knowledge and Justification*. New York: Oxford University Press.
- _____. 2015. "Four Challenges to the A Priori – A Posteriori Distinction." *Synthese* 192: 2701-24.
- Jeshion, Robin. 2011. "Experience as a Natural Kind: Reflections on Albert Casullo's *A Priori Justification*." In *What Place for the A Priori?*, eds. M. J. Shaffer and M. L. Veber. Chicago: Open Court.

Kitcher, Philip. 1983. *The Nature of Mathematical Knowledge*. New York: Oxford University Press.

_____. 2000. "A Priori Knowledge Revisited." In *New Essays on the A Priori*, eds. P. Boghossian and C. Peacocke. Oxford: Oxford University Press.

Moser, Paul. 1989. *Knowledge and Evidence*. Cambridge: Cambridge University Press.

Williamson, Timothy. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell Publishing.

_____. 2013. "How Deep is the Distinction between A Priori and A Posteriori Knowledge?" In *The A Priori in Philosophy*, eds. A. Casullo and J. Thurow. Oxford: Oxford University Press.

Notes

1. See Casullo (2013) for a broader survey of the challenges.
2. See Casullo (2012b) for a discussion of this argument.
3. All page references in the text are to Williamson (2013) unless otherwise indicated.
4. Williamson's dilemma is puzzling. His goal is to show that the a priori – a posteriori distinction is superficial or insignificant. With respect to the bottom-up approach, he maintains that it runs the risk of introducing a distinction of little significance. With respect to the top-down approach, however, he maintains that it runs the risk that the concept of the a priori is vacuous. Here it is important to distinguish two different challenges to the concept of the a priori: (1) the concept is vacuous, and (2) the concept is insignificant. The first challenge is familiar. It is the challenge of radical empiricists who deny the existence of a priori knowledge. Williamson's goal, however, is to introduce a novel challenge, one that does not deny the existence of a priori knowledge. Hence, Williamson's initial presentation of his argument does not even purport to show that the top-down approach introduces a concept of the a priori that is insignificant.

One might suggest that if a concept is vacuous then it is also insignificant. This suggestion is mistaken. The question, "Does the Higgs boson exist?" is theoretically significant even if the answer to it should turn out to be negative. On the other hand, the question, "Do green tulips exist?" is theoretically insignificant even if the answer is

affirmative. So vacuity and significance do not go hand in hand. See Casullo (2013) for a taxonomy of challenges to the a priori – a posteriori distinction.

5. I will address whether (K1) and (K2) can be defended by top-down standards in section 3.
6. A referee suggests that Williamson might maintain that Norman's knowledge of (1) is a paradigm case of a priori knowledge and that Norman's knowledge of (2) is a paradigm case of a posteriori knowledge. I don't find such appeals to paradigm cases compelling. It is plausible to maintain that knowledge of (1) is a paradigm of a priori knowledge in the sense that a theory of a priori knowledge on which knowledge of (1) is not a priori is prima facie deficient. It must offer some theoretical justification for this deviation from our pre-theoretical classification of the cases. The same is true of a theory of a posteriori knowledge on which knowledge of (2) is not a posteriori. This appeal to paradigm cases, however, does not support the stronger claim that *any* way of coming to know (1) is a paradigm case of a priori knowledge. Presumably, one can know (1) on the basis of testimony, but it does not follow that testimonial knowledge of (1) is a paradigm case of a priori knowledge.
7. Williamson has available a strategy for answering this objection. He could maintain that when the a priori – a posteriori distinction is introduced bottom-up by examples, the examples of a priori ways of knowing include knowing (1) (or relevantly similar propositions) via steps (A1) - (A2) (or relevantly similar steps) and the examples of a posteriori ways of knowing include knowing (2) (or relevantly similar propositions) via steps (B1) - (B2) (or relevantly similar steps). This strategy, however, does not strike me as very promising since, as Williamson points out, the distinction is primarily a classification of specific ways of knowing and it would be pointless, to say the least, to introduce a distinction among ways of knowing by means of examples that do not differ even superficially.
8. The issue of overdetermination looms large in the literature on a priori knowledge. For example, I contend (in Casullo 2005) that the arguments of Mill and Quine against the existence of a priori knowledge fail because they overlook the possibility of epistemic overdetermination. I also contend (in Casullo 2012b) that Williamson's (2007) argument in support of (1.1) runs into difficulty because it overlooks the possibility of epistemic overdetermination.
9. A referee wonders whether the cognitive processes involved in traditional rationalist accounts of a priori knowledge differ substantially from the cognitive processes involved in Williamson's account. One clear difference is that the traditional rationalist accounts invoke a grasp of universals such as *red* and *crimson*, together with insight into their (relational) properties such as that *red* includes *crimson*. Since these cognitive processes are directed toward abstract entities, proponents of the view maintain that they are not perceptual.

10. See, for example, Bealer (1992) and BonJour (1998).
11. A referee astutely notes that many epistemologists contend that we can rightly take our basic belief-forming processes to be reliable without relying on inferential arguments for their reliability. So, the referee suggests, perhaps Williamson can maintain that we have reason to trust that the process of imagination involved in the Norman cases is reliable (since it is a basic process) in the absence of reason to believe that it is unreliable. As I understand this contention regarding basic processes, it states that if S's belief that p is based on (produced by) a basic process, then S's belief that p is *prima facie* justified. Putting the point somewhat differently, the contention states that the deliverances of basic processes are basic beliefs—i.e., they are *prima facie* justified and their justification does not derive from other justified beliefs. But it does not follow from this contention that the belief that some basic process is reliable is also basic. My claim is not that, in order to have knowledge on the basis of an exercise of the process of imagination involved in the Norman cases, one must first have evidence that the process is reliable. I am claiming only that Williamson's contention that the process is sufficiently reliable to bestow positive epistemic status on the beliefs that it produces requires supporting evidence.
12. A referee rejects this claim on the grounds that it could turn out that, when we try to systematize the distinction introduced bottom-up, we cannot do so in any plausibly coherent way; in which case, we can conclude that there is no *a priori* knowledge. This argument is questionable. It is important to distinguish two different challenges to the *a priori* – *a posteriori* distinction: the distinction is superficial (but coherent) and the distinction is incoherent. The scenario envisaged by the referee provides evidence that the distinction is incoherent. But, if the distinction is incoherent, it does not follow that there is no *a priori* knowledge. If the distinction is incoherent, then the claim that there is no *a priori* knowledge is also incoherent. See note 4 and Casullo (2013) for further discussion of the difference between the two challenges.
13. Williamson also maintains that the top-down approach faces the risk that there is no *a priori* knowledge. The fact that the top-down approach faces this risk is a mark in its favor. It shows that the top-down approach to introducing the *a priori* – *a posteriori* distinction leaves open the primary question that it was introduced to address.
14. One might suggest, in defense of Williamson, that our puzzlement when we consider the two Norman cases has an alternative source:

(P3) When we consider the two Norman cases, we are puzzled as to whether the term 'independent' applies.

(P3) is also false. There is controversy over how to understand the term 'independent'. That controversy revolves around a strong and weak reading of the term:

(SI) The justification of S's belief that p is independent of experience if and

only if it does not derive from experience and cannot be defeated by experience.

- (WI) The justification of S's belief that p is independent of experience if and only if it does not derive from experience.

The two Norman cases, however, do not involve any defeating conditions. Consequently, it is not unclear or indeterminate whether 'independent' applies to them as it might be in cases where defeaters for justification are involved.

15. Thanks to Elia Zardini for pressing me to clarify the argument of this paragraph.
16. I have stated (JE) somewhat vaguely to leave open a number of different views that appear to entail it, such as (a) the concept of justification is analyzable in terms of the concept of evidence, (b) the concept of justification and the concept of evidence are coextensive, (c) one's justification supervenes on one's evidence, and (d) only evidence justifies.
17. See Moser (1989) for such a theory.
18. Thanks to Elia Zardini for alerting me to a problem in my original formulation of this claim.
19. I offer a response to the challenge in Casullo (2003). For subsequent discussion of that response, see Brueckner (2011), Jeshion (2011) and Casullo (2011).
20. Thanks to Elia Zardini and to an anonymous referee for their careful reading of my paper and for their insightful comments.