The Theoretical Significance of the A Priori/A Posteriori Distinction[*]

Joshua Schechter

1. Introduction

The distinction between a priori and a posteriori knowledge is a familiar one. It is straightforward to generate a long list of examples of (apparent) a priori knowledge and a long list of examples of (apparent) a posteriori knowledge. For instance, here is a list of claims that I presumably know a priori:

- *Logical truths*. Every walrus is a walrus.
- *Metalogical truths*. Conjunctions logically entail their conjuncts.
- *Mathematical truths*. 5+7=12.
- *Simple conceptual truths*.[1] Every bachelor is unmarried.
- *Judgments about cases*. In [such-and-such Gettier case], Smith does not know.
- *Color relations*. No shade of red is a shade of green.
- *Modal truths*. It is metaphysically possible for there to be a bicycle that is a hundred feet tall.
- *Moral principles*. It is morally wrong to cause gratuitous suffering.
- *Epistemic principles*. Typically, someone who knows a conjunction is in a position to know each of its conjuncts.
- *"Superficially" contingent a priori truths*.[2] All actual philosophers are philosophers.
- *"Deeply" contingent a priori truths*.[3] My cognitive and perceptual faculties are broadly reliable.

---

[1] This label is not meant to rule out the view that other items on the list are conceptual truths, too.

[2] The distinction between superficially and deeply contingent truths is due to Evans (1979). A true claim counts as superficially contingent if there is a possible world in which it is false. A true claim counts as deeply contingent if the laws of semantics do not guarantee that it is actually true. On Evan's definitions, the deeply contingent truths are a subset of the superficially contingent truths. The example of a superficially contingent a priori truth above is meant to be one that is not also deeply contingent. Evans claimed that while there is "nothing particularly perplexing" about the existence of a superficially contingent a priori truth, it would be "intolerable" for there to be a deeply contingent a priori truth. (Evans 1979, pp. 161-2).

[3] This is the most controversial item on the list. The idea here comes from discussions of skepticism. It is plausible that whenever a perceptual faculty yields knowledge (or justification), one is in a position to know (or justifiably believe) that the faculty is reliable. What explains this correlation? There are at least three possibilities. First, it could be that for a perceptual faculty to yield knowledge requires one to antecedently be in a position to know that the faculty is reliable. Presumably, on pain of regress or vicious circularity, the knowledge that one is antecedently in a position to have (i.e., that the faculty is reliable) must be a priori. See Wright (2002) and White (2006) for views in this general ballpark. Second, it could be that the direction of explanation goes in the other direction. For instance, perhaps one can acquire knowledge that a perceptual faculty is reliable by using that faculty in an "offline" or suppositional way. Assuming that offline uses of perceptual faculties yield a priori knowledge, on this view we are again in a position to have a priori knowledge that the faculty is reliable. See Cohen (2010) and Wedgwood (2013) for views in this ballpark. Finally, it could be that there is a common factor that explains why the perceptual faculty

And here is a list of claims that I presumably know a posteriori:

- *Directly observational truths*. It is currently raining.
- *Truths about the past*. Dinosaurs used to roam the earth.
- *Truths about the future*. The water in the pot on the stove will soon come to a boil.
- *Contingent generalizations*. No skyscraper is more than one mile tall.
- *Statistical claims*. The global literacy rate is about 86%.
- *Truths about linguistic meaning*. The word "table" stands for a kind of furniture.
- *Constitutive claims*. Water is composed of H_2O molecules.[4]
- *Laws of nature and scientific theories*. General relativity well describes the behavior of highly massive objects.

It is notable that each of these lists is very heterogeneous. There seems little in common between knowing that every walrus is a walrus, knowing that it is morally wrong to cause gratuitous suffering, and knowing that one's cognitive and perceptual faculties are broadly reliable. These three claims have different subject matters, and the ways in which we come to know them are not obviously similar. And the same holds true for the examples of a posteriori knowledge. There seems little in common between knowing that it is currently raining, knowing that the word "table" stands for a kind of furniture, and knowing that general relativity well describes the behavior of highly massive objects. Again, these three claims have different subject matters, and, again, the ways in which we come to know them are not obviously similar.

This provides some (perhaps weak) evidence that the a priori/a posteriori distinction is not a very natural one in the sense that it does not "carve nature at its joints".[5] The distinction may resemble the distinction between kinds of things that have a name in English starting with the letter "k" and kinds of things that don't have a name in English starting with the letter "k" – that is to say, a legitimate distinction but one that doesn't have any real significance.

In recent years, several philosophers have argued for this very conclusion – that the a priori/a posteriori distinction is legitimate but non-natural and theoretically unimportant. For instance, John Hawthorne writes, "my own externalist commitments … lead me to think that the a priori–a posteriori distinction is not a particularly natural one, and hence that its importance to epistemology has been grossly overestimated."[6] Similarly, Timothy Williamson writes, "although a distinction between a priori and a posteriori knowledge (or justification) can be drawn, it is a superficial one, of little theoretical significance."[7]

---

yields knowledge and why we are in a position to know that the faculty is reliable. Presumably, the common factor does not involve experience. So, on this view, too, it is plausible that this knowledge that the faculty is reliable will be a priori. See Silins (2008) for discussion of views in this ballpark. (The case of cognitive faculties is analogous.)

[4] This is one of Kripke's examples of an a posteriori contingent truth (although Kripke uses the identity "Water is H_2O" rather than a claim about constitution). See Kripke (1980, lecture 3).

[5] See Quine (1969) and Lewis (1983) for this usage of "natural". This usage of "natural" is not the one at issue in discussion of whether we should be naturalists about some domain.

[6] Hawthorne (2007, p. 201).

[7] This quote is taken from the abstract to Williamson (2013). It does not appear in the main body of the article, but the article contains several passages making the same claim.

I have some sympathy with the claim that the distinction between a priori and a posteriori knowledge, at least as traditionally understood, does not capture a natural or theoretically useful distinction. But Hawthorne and Williamson also seem to think that there is no natural or theoretically useful distinction in the vicinity, and that seems too strong a conclusion to draw. My own view is that there is a theoretically useful distinction – or better, distinctions – in the ballpark of the traditional distinction.[8] Giving up on anything like an a priori/a posteriori distinction is throwing out the baby with the bathwater.

In this paper, I do two main things. First, I respond to the most prominent recent challenge to the significance of the a priori/a posteriori distinction – the central argument in Williamson (2013).[9] Second, I discuss the question of what the theoretical significance of the a priori/a posteriori distinction is.

This paper will proceed as follows. In the next section, I present the a priori/a posteriori distinction as it is typically developed. I then turn to Williamson's challenge to the significance of the distinction. Williamson points out that we often use the same cognitive mechanisms in coming to have a priori and a posteriori knowledge. So how could it be, asks Williamson, that there is a "deep epistemological difference" between the two? In section 3, I present this challenge and in section 4, I respond to it. I argue that there is an important disanalogy between Williamson's central example of a case of a priori knowledge and his central example of a case of a posteriori knowledge. Although the beliefs in the two cases are formed in similar ways, the ways in which their justification can be defeated are different. This suggests that there is an important epistemological difference between the two cases, one that cannot be captured in terms of the cognitive mechanisms used to form the beliefs.

Although Williamson's argument is unsuccessful, there remains the question of just what the theoretical significance of the a priori/a posteriori distinction is. In section 5, I turn to that question. I argue that the point of the distinction is not to enable us to represent some joint in nature, but rather to help us to identify epistemological problem cases. We understand – more-or-less – the epistemology of simple perceptual knowledge. The epistemology of non-perceptual knowledge is far less clear. The purpose of labeling a case of knowledge as a priori is to claim that its epistemology should not be assimilated to the epistemology of perception. Instead, it is something of a puzzle case.

This proposal has an important implication, which I discuss in section 6. There are several ways in which a case of knowledge can be disanalogous to a simple case of perceptual knowledge. Two disanalogies are perhaps the most important: the explanation of how it is the belief is justified may be different, and the explanation of how it is the thinker has a true belief may be different. It is natural to think that in a simple case of perceptual knowledge, the explanation of how the belief is justified centrally involves phenomenality, and the explanation of how the thinker came to have a true belief centrally involves a causal relation between the belief and what

[8] Peacocke (2004, ch. 6.3) identifies two species of the a priori, the "judgmentally a priori" and the "contentually a priori". His distinctions are conceptually closer to one another than the distinctions that I will ultimately argue for below.
[9] An earlier argument, relying on similar considerations, appears in Williamson (2007, ch. 5.5).

the belief is about. If so, the two disanalogies come to the following: (i) the justification of the belief does not involve phenomenality, and (ii) the belief does not stand in a causal relation to what the belief is about. When beliefs about some subject matter fit either (i) or (ii), an epistemological puzzle arises. How can our beliefs about some subject matter be justified if not ultimately in terms of how things seem to us? How is it that we get things right about some subject matter if our beliefs do not stand in a causal relation to what they are about? So there is more than one kind of epistemological puzzle to solve. This suggests that there is an important theoretical role for (at least) two distinctions in the ballpark of the traditional a priori/a posteriori distinction.

2. The A Priori and A Posteriori

In the previous section, the a priori/a posteriori distinction was introduced using examples. But it is more common (and more satisfying) to introduce it using definitions. And there are familiar definitions to provide:[10] A thinker has a priori knowledge that p just in case the thinker knows that p and the thinker's knowledge that p is independent of experience. A thinker has a posteriori knowledge that p just in case the thinker knows that p and the thinker's knowledge that p depends on experience.[11,12]

This is not the original understanding of the terms "a priori" and "a posteriori". These terms were earlier used to mark a different distinction. An a priori demonstration was a demonstration from the cause (or, better, explanatory ground) of the fact to be proved. An a posteriori demonstration was a demonstration from an effect of the fact to be proved. This was the prevailing usage of the terms from the 14th through the late 17th centuries.[13] The contemporary usage of "a priori", as

---

[10] There are other uses of "a priori" in non-philosophical speech and writing. I've heard the term used to mean something like "on the basis of general considerations rather than specific evidence", "on the basis of a back-of-the-envelope calculation rather than a detailed derivation", and "on the basis of our current evidence and not any further empirical investigation". I put such uses of the term aside here. Presumably, they trace back to the philosophical use(s) of the term.

[11] There is a small refinement that should be made to these definitions. A thinker may have two sets of grounds for believing a claim, one this is independent of experience and one that depends on experience. The thinker may believe the claim on both sets of grounds. So long as the first grounds epistemically suffice for knowledge, we should count the thinker as having a priori knowledge of the claim. If the second set of ground also suffice for knowledge, we should count the thinker as having a posteriori knowledge of the claim, as well. I'll leave this refinement implicit in what follows.

[12] These definitions concern the application of "a priori" and "a posteriori" to knowledge. The terms are also commonly applied to the epistemic justification of beliefs, with analogous definitions. It is also common to talk about whether propositions are a priori – a proposition counts as a priori just in case it can be a priori known (or justified). Finally, the terms are occasionally applied to entire subject matters – a subject matter counts as a priori just in case the propositions it concerns are (at least centrally) a priori propositions and as a posteriori otherwise.

[13] For example, an early characterization of "a priori" and "a posteriori" appears in William of Ockham's *Summa Logicae* [c. 1323] III.2, ch. 17. Ockham uses the terms to mark the distinction between demonstrations whose premises are explanatorily prior to the conclusion and demonstrations whose premises are not explanatorily prior to the conclusion. Ockham identifies the a priori/a posteriori distinction with the earlier distinction between demonstrations *propter quid* and demonstrations *quia*, which is the Latin translation of Aristotle's distinction in the *Posterior Analytics* I.13 between demonstrations "of the reasoned fact" (διότι) and demonstrations "of the fact" (ὅτι). (Indeed, *Summa Logicae* III.2 is essentially a commentary on the *Posterior Analytics*.)

meaning something like "independent from experience", emerged in the 18ᵗʰ century and was popularized largely due to the influence of Immanuel Kant.[14]

Despite the fact that the terms "a priori" and "a posteriori" only gained this usage relatively late, the idea that there is an important distinction between knowledge that depends on experience and knowledge that is independent of experience goes back much further. As Alberto Coffa writes, "One of the basic intuitions behind almost every epistemology since Plato's is that there are two radically different types of claims: the a priori and the rest."[15] Indeed, it is difficult to understand what Plato's theory of recollection is meant to do, if it is not an answer to the question of how thinkers can have a priori knowledge of subjects like ethics and mathematics.[16]

In the late 17ᵗʰ century, essentially the same usage can be found, for example, in Arnauld and Nicole's *Port Royal Logic* [1662/1683], part IV, ch. 1. A remnant of this usage can be found as late as Frege (1974) [1884], sec. 3, where a truth is defined to be a priori if its ultimate explanatory grounds are general laws not themselves in need of proof. For discussions of the history of the a priori/a posteriori distinction, see Schepers et al. (1971) and Mittelstrass (1977).

[14] Kant (1998, A 2/B 2-3) [1781/1787]. But see Smit (2009), for an argument that Kant was also operating with the earlier – from the explanatory ground – definition of "a priori".

Notice that there are two main changes from the earlier definition of the a priori to the contemporary one. First, the definition doesn't apply to demonstrations but to knowledge (or justified belief, etc.). Second, the definition no longer makes use of the notion of explanatory priority but the notion of dependence on experience. The first change is understandable: Given a distinction between kinds of demonstrations, it is natural to apply the distinction to the conclusions of demonstrations, and perhaps to then extend its application still further. The second change is more puzzling. Adams (1994, pp. 109-10) argues that Leibniz used the terms "a priori" and "a posteriori" in the earlier sense, but that it followed from his substantive philosophical views that a priori knowledge coincides with knowledge that is non-empirical. Adams thus suggests that Leibniz therefore may have "played a crucial role in the transformation of the meaning of 'a priori'" (p. 110).

Kant was not the first to define "a priori" and "a posteriori" in terms of experience. Related characterizations can be found in Wolff (1728, section 663), Baumgarten (1761, section 316), and Lambert (1764, sections 534-45). Wolff's characterizations are interestingly different from the contemporary ones. On his characterization, roughly, knowledge is a priori if it was established as true in part using inference from known propositions. Knowledge is a posteriori if it was established as true using only sense experience. So on Wolff's understanding of the terms, a claim that is known using a combination of experience and reason would count as a priori. By contrast, on the contemporary understanding, such a mixed case would count as a posteriori. The characterizations in Baumgarten and Lambert are closer to the contemporary one. It should be mentioned that in his earlier *Metaphysics* (2013, section 24) [1739], Baumgarten uses the terms "a priori" and "a posteriori" in a way that suggests that he was then using the earlier – from the explanatory ground – understanding. So his usage of the terms seems to have shifted over time. (Hume (2007, 4.1) [1748] uses the term "a priori" but doesn't provide an explicit characterization of it. He appears to be using the term in its earlier sense, but he does contrast knowledge attained by a priori reasoning with knowledge attained through experience. Hume's distinction between relations of ideas and matters of fact is close to the contemporary a priori/a posteriori distinction but is perhaps better assimilated to the analytic/synthetic distinction.)

[15] Coffa (1991, pp. 1-2). Coffa doesn't tell us which part of Plato he has in mind. But a natural suggestion is Plato's discussion of the differences between our epistemic access to the "visibles" (ordinary objects, shadows, and reflections) and our epistemic access to the "intelligibles" (the Forms and mathematical objects) in *Republic*, Book 6.

[16] See Plato's *Meno* and *Phaedo* for the theory of recollection, which is (roughly) the theory that our souls had a kind of direct acquaintance with the Forms and the mathematical objects before our birth, and our knowledge of morality, mathematics, and other a priori domains is due to recollection of this earlier acquaintance.

The conception of a priori knowledge as knowledge that is independent of experience requires precisification. More needs to be said on how we should understand "experience" and how we should understand "independent of". There are challenges facing both tasks.[17]

First consider the question of how we should understand "experience", at least for the purpose of defining the a priori. One issue here is whether we should count only experiences of the external world, or whether we should also include experiences that are purely internal, such as introspective and proprioceptive experiences. A second, perhaps deeper, issue concerns the question of what makes a mental state count as an experience. Is it that the mental state has a phenomenal aspect – that there is something it is like to be in that mental state?[18] Or is it that the mental state satisfies some kind of causal constraint – for instance, the mental state was caused by what it is about?[19] Or is it that experiences form some kind of natural kind – a mental state is an experience if it is the same kind of psychological or neurophysiological state as a paradigmatic example of an experience?[20] Or should we endorse a hybrid account, or perhaps something else entirely? An answer to the question of what makes a mental state an experience is needed so that we can delimit the boundary of the a priori and so we can better understand what the a priori is supposed to be.

There are difficulties facing the various candidate answers. Against a phenomenal conception of experience, there is a problem with sufficiency – if we allow that thinkers can have a priori knowledge via rational insights, understood to be phenomenally-laden cognitive states, we should not count such phenomenal states as experiences. There may also be a problem with necessity – if it is possible for someone with blindsight to know facts about their surroundings directly via blindsight, such knowledge should presumably not count as a priori.[21] Against a causal conception, presumably hallucinations should count as experiences even though they are almost never caused by what they are about. And against a natural kind conception, it seems conceptually possible for there to be creatures that have psychologies and neurophysiologies that are very different from our own, yet nevertheless have experiences. (Think space aliens, or even octopuses.) If their psychological and neurophysiological architectures are sufficiently different from our own, their experiences will not share the same psychological or neurobiological kinds as our experiences.[22]

---

[17] See BonJour (1998, pp. 7-11).

[18] Block (1995) identifies experiences as mental states with phenomenal properties. He also makes use of the "what it is like" terminology originally due to Nagel (1974).

[19] BonJour (1998, p. 8) and McGinn (1975/1976, p. 198) propose causal accounts of experience. Presumably, there needs to be a restriction to causes that are in some sense direct or immediate. Otherwise, if I were to see smoke and come to believe that there is a fire, my belief would count as an experience of the fire since the fire did (mediately) cause my belief that there is fire.

[20] See Casullo (2003, sec. 6.2) for discussion of the options. Casullo himself favors the natural kind approach.

[21] In actual cases of blindsight, as described in the empirical literature, subjects do not spontaneously form beliefs about their environment but are prompted to guess. So actual subjects with blindsight presumably do not gain knowledge about their environment directly via blindsight. But we can imagine subjects with a souped-up kind of blindsight who spontaneously, noninferentially, and reliably form beliefs about their environment. This is close to what Block (1995) calls "super-blindsight".

[22] This is a generalization of the case of "Martian pain" in Lewis (1980).

These difficulties notwithstanding, there does seem to be something right in both the phenomenal and causal approaches. It seems to be an important fact about experience that experiences have phenomenal properties – that there is something it is like to have an experience. It also seems to be an important fact about experience that experiences, at least when they are non-defective, stand in a causal relation to what they are about. Perhaps the most promising avenue for formulating a single general account of experience is to adopt a hybrid approach that combines elements of both the phenomenal and the causal conceptions, and which avoids some of the difficulties by appealing to teleological or functional notions. For example, one might say that a mental state is an experience if it was formed by a psychological process that has the function of both (i) making use of causal relations to acquire information about some part of the world and (ii) providing that information to the thinker via phenomenality.[23] Such a proposal avoids many of the difficulties listed above. But, to anticipate some of the discussion below, I don't see any need to formulate a single general account of experience. There is more than one legitimate notion of experience that can be used in formulating a definition of the "a priori", and there is more than one definition that has an important theoretical role to play.

Putting that issue aside for now, let me turn to the question of how to understand "independent of". There is a familiar difficulty here, and a familiar response. Consider the fact that no shade of green is a shade of red. It is natural to think that our knowledge of this fact is a priori. But it is also natural to think that this knowledge is not completely independent of experience. A person with normal vision will typically possess the concept *green* on the basis of having had visual experiences of green things, and so their knowledge that no shade of green is a shade of red will depend on their having had various experiences.[24] That is the familiar difficulty with the definition.

The familiar response is to distinguish two ways in which experience may be relevant to a piece of knowledge. It may play a "merely enabling" role or it may play an "evidential" role. In the case of knowledge that no shade of green is a shade of red, the role of green experience is merely to enable us to possess the concept *green*. It is not to provide evidence for the claim that no shade of green is a shade of red.[25]

The idea, then, is that experience plays a merely enabling role in knowledge of a proposition if it does not provide evidence for the truth of the proposition but (for instance) is only necessary for possessing one of the ingredient concepts of the proposition. An experience plays an evidential role in knowledge of a proposition if it genuinely provides evidence for the truth of the proposition. A priori knowledge is compatible with experience playing a merely enabling role.

---

[23] Casullo (2011, p. 131) extracts (and rejects) a proposal like this from the discussion in Jeshion (2011).

[24] A congenitally blind person may possess the concept *green* in some other way, but presumably will possess the concept in a way that also depends on experience – for instance, experience of testimony.

[25] The basic idea behind this response appears in Kant's *Critique of Pure Reason* B 3. I borrow the terms "enabling" and "evidential" from Williamson (2013). I don't know who is originally responsible for this terminology. Burge (1993) makes use of a distinction between "enabling" and "justifying or entitling". The enabling/evidential distinction is broadly analogous to the distinction in causal explanation between causally enabling conditions and genuine causes.

But it is not compatible with experience playing an evidential role. That is why our knowledge that no shade of green is a shade of red counts as a priori.

This familiar response ought to be complicated a bit to handle a few subtle issues. First, we should distinguish possessing a concept in the sense of having competence with the concept from being able to think thoughts that contain the concept as a constituent. These are not the same. It is easy to think thoughts that contain a specific concept as a constituent despite lacking any real competence with the concept. For instance, suppose I overhear someone talking about Galois theory and I have no idea what Galois theory is. That is enough for me to be able to think thoughts about Galois theory – I can wonder what Galois theory is and whether it would be useful for me to learn it. For the purpose of characterizing a priori knowledge, any experience that is required to possess a concept in the sense of having competence with the concept should be classified as playing an enabling role. So, playing an enabling role can include more than just enabling one to be able to think thoughts containing the concept as a constituent.

Second, it could be that there are cases of a priori knowledge in which experience enables one to possess a concept that is not a constituent of the proposition that is known, but is a constituent of a proposition that appears in the argument or justification for the known proposition. There are many examples of mathematical proofs in which notions appear in the proof that do not appear in the claim to be proved. (The proof of Fermat's Last Theorem is an extreme example of this.) If experience plays a role in enabling one to possess such a concept, that is compatible with the knowledge of the proved claim counting as a priori.

Third, to play a so-called "evidential role" in our knowledge of a claim, experience presumably need not provide direct evidence for the claim. It would suffice for experience to defeat counterevidence against the claim. It would also suffice for experience to provide a non-evidential warrant for the claim (if experience can play such a role).[26] What is crucial for a posteriori knowledge is not that experience provides evidence, but that it somehow contributes to the epistemic warrant for the claim.[27]

These complexities aside, the enabling/evidential distinction seems tolerably clear, and appealing to this distinction seems to answer the concern that paradigmatic cases of a priori knowledge (such as our knowledge that no shade of green is a shade of red) can depend on experience in some way. However, Williamson's argument against the naturalness and theoretical significance of the a priori/a posteriori distinction directly takes aim at the use of the enabling/evidential distinction in drawing the contrast. So let me now turn to Williamson's argument.

---

[26] It is plausible that there can be non-evidential kinds of warrant. For instance, what is my evidence for the claim that 2+3=5 or the claim that bachelors are unmarried? Assuming we have a priori warrant for the claim that my perceptual faculties are reliable or the claim that the world contains projectible regularities, how is that warrant based on evidence? For simplicity, in this paper, I treat "justification" and "warrant" as synonyms, to encompass both evidential and non-evidential kinds of warrant.

[27] For this reason, Boghossian (2002, p. 139) contrasts enabling roles with epistemic (rather than evidential) roles. That is an improvement, but in what follows I'll stick with Williamson's terminology.

3. Williamson's Argument

Williamson grants that there is a legitimate distinction between a priori and a posteriori knowledge. But he argues that the distinction "does not cut at the epistemological joints" and that it is "unhelpful for theoretical purposes".[28] (Notice that these are two distinct conclusions, though Williamson seems to identify them.)

The structure of Williamson's argument is straightforward. He describes (what he takes to be) a clear case of a priori knowledge and (what he takes to be) a clear case of a posteriori knowledge. He argues that the differences between the two cases are epistemologically superficial, and therefore the distinction between the a priori and the a posteriori does not carve nature at its joints.

This kind of argument can be powerful. If there are clear cases of X and clear cases of Y which differ only superficially, in a way that makes it difficult to see how the difference could have any real significance, this would seem to cast doubt on the naturalness of the distinction between the Xs and the Ys.[29] Indeed, in previous work, I have argued in just this way for the (very un-Williamsonian) claim that the distinction between knowledge and non-knowledge does not carve at the epistemological joints – there are clear cases of knowledge that closely resemble clear cases of non-knowledge, and it is difficult to see how anything of epistemic significance could turn on the subtle differences between the two kinds of cases.[30] So if there is a problem facing Williamson's argument, it does not concern its overall strategy but the details of its implementation.

Williamson chooses as his central examples a case of knowing

> (1) All crimson things are red

and a case of knowing

> (2) All recent volumes of *Who's Who* are red.

Williamson claims that in normal cases, knowledge of (1) is a priori and knowledge of (2) is a posteriori. (I think he is clearly right about this.) He then introduces a subject, Norman, and describes how Norman comes to know (1) and (2) in some detail:

> Suppose that Norman acquires the words 'crimson' and 'red' independently of each other, by ostensive means. He learns 'crimson' by being shown samples to which it applies and samples to which it does not apply, and told which are which. He learns 'red' in a parallel but causally independent way. … Through practice and feedback, he

---

[28] Williamson (2013, p. 295).

[29] That a clear case of X differs only superficially from a clear case of Y does not *entail* that the distinction between the Xs and Ys is not a natural one. But it does provide support for that conclusion.

[30] See Schechter (2017, p. 138-9). This raises the question of how the a priori/a priori distinction could carve at the joints if the knowledge/non-knowledge distinction does not. As will become clear below, my view is that the a priori/a posteriori distinction is not fundamentally a distinction concerning knowledge. Rather, there are at least two distinctions, one concerning true belief and the other concerning justification. I'm also not convinced that the a priori/a posteriori distinction does carve at the joints – it is theoretically useful, but that's a different matter.

becomes very skilful in judging by eye whether something is crimson, and whether something is red. Now Norman is asked whether (1) holds. … [H]e assents to (1) after brief reflection on the colours crimson and red, along something like the following lines. First, Norman uses his skill in making visual judgments with 'crimson' to visually imagine a sample of crimson. Then he uses his skill in making visual judgments with 'red' to judge, within the imaginative supposition, 'It is red'. This involves a general human capacity to transpose 'online' cognitive skills originally developed in perception into corresponding 'offline' cognitive skills subsequently applied in imagination. … As a result of the process, Norman accepts (1). Since his performance was sufficiently skilful, background conditions were normal, and so on, he thereby comes to know (1). Williamson (2013, p. 295)

[Norman] learns the complex phrase 'recent volumes of *Who's Who*' by learning 'recent', 'volume', '*Who's Who*' and so on. … 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. … [H]e 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. … 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). Williamson (2013, p. 296)

Williamson goes on to argue that although Norman's knowledge of (1) is a priori and his knowledge of (2) is a posteriori, the cases are very similar. He writes:

The cognitive process 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 shallow. Williamson (2013, pp. 296-7)

In further support of this conclusion, Williamson considers the prospect of appealing to the distinction between enabling and evidential roles of experience to distinguish between Norman's knowledge of (1) and his knowledge of (2). Williamson argues that in neither case is the role of experience merely enabling or straightforwardly evidential. Rather, experience plays an intermediate role. That is because the role of experience in both cases is not merely to enable Norman to possess the concepts *crimson*, *red*, and *recent volume of Who's Who*.[31] Someone – Williamson calls him "Norbert"[32] – could acquire competence with the concepts *crimson* and *red* by ostensive means (just like Norman), but have less practice with them than Norman and be less

---

[31] Jenkins and Kasaki (2015, p. 2727) point out that Williamson's discussion moves back and forth between claims about words and claims about concepts, but I think this ultimately turns out to be harmless.
[32] Williamson (2013, pp. 297-8).

skillful in imagining samples of crimson. Norbert may count as competent with *crimson* and *red* but be unable to come to know that all crimson things are red via imagining a crimson object. If that's right, then the role of experience in Norman's case goes beyond what is required for Norman to possess the concepts *crimson* and *red*. So, Williamson argues, Norman's experience does not play a merely enabling role. But it does not play an evidential role either, either, since the role of experience is to hone Norman's skills of imagining and identifying colors and not to provide evidence for color relationships.

As stated, this argument is contestable. It seems to rely on the claim that the only way for experience to play a merely enabling role is for it to contribute to concept possession. That is an overly narrow conception of when experience plays a merely enabling role. For instance, in working out a long and complex mathematical proof, one may rely upon experience (e.g., in making use of a chalkboard). This experience need not contribute to possessing a concept. But such experience ought to be counted as playing a merely enabling rather than a genuinely evidential role.

Williamson does, however, have a better argument against appealing to the distinction between enabling and evidential roles of experience. He notes that the role of experience seems to be almost exactly the same in the case of Norman's knowledge of (1) and in the case of Norman's knowledge of (2). Given this, it is difficult to see how to draw a principled distinction so that experience plays a merely enabling role in Norman's knowledge of (1) and an evidential role in Norman's knowledge of (2). So, Williamson argues, the distinction between merely enabling and genuinely evidential roles can't be used to distinguish between Norman's knowledge of (1) and (2).

There are several potential responses to this argument. One response would be to grant that Williamson is right that there are only minor differences between how Norman comes to know (1) and how Norman comes to know (2), but to claim that Norman's knowledge of (1) and (2) fall on the same side of the a priori/a posteriori divide, and further, that there are genuine cases on the other side of the divide. On the most plausible way of fleshing out this view, Norman's knowledge of (1) and his knowledge of (2) both count as a posteriori, but there are other ways of coming to know (1) that are genuinely a priori.[33] (Or perhaps knowledge of (1) is always a posteriori, but we can have genuine a priori knowledge of other propositions, such as simple logical and mathematical claims.)

Williamson is sensitive to this kind of response. He argues – convincingly to my mind – that Norman's way of coming to know that all crimson things are red is not highly idiosyncratic. We often come to have knowledge of (apparently) a priori claims using the imagination, where the skills we deploy in the imagination were honed by perceptual experience. Indeed, Williamson argues that we use the imagination in this way even in simple cases of logical and mathematical knowledge, such as our knowledge of the reflexivity of identity and of the powerset axiom in set

---

[33] Casullo (this volume) defends this response to Williamson.

theory.[34] But even if Williamson is wrong about this, it is intuitively correct to count Norman's knowledge of (1) as a priori and his knowledge of (2) as a posteriori. So to respond to Williamson by lumping together Norman's knowledge of (1) and (2) is to pursue the wrong avenue.

A better response is to claim that there is, after all, a crucial disanalogy between Norman's knowledge of (1) and Norman's knowledge of (2). One way to try to develop such a disanalogy is to consider what contributes to conceptual mastery (rather than mere conceptual competence). The idea is that becoming more skillful in imagining and identifying colors based on their visual appearance contributes to gaining greater mastery of color concepts. By contrast, becoming more skillful in imagining and identifying recent volumes of *Who's Who* based on their visual appearance does not contribute to gaining greater mastery of the concept *recent volume of Who's Who*. So the role of experience in Norman's knowledge of (1) is disanalogous to the role of experience in Norman's knowledge of (2). On this approach, experience plays a merely enabling role in a thinker's knowledge of some proposition if (for example) the experience merely provides the thinker with some amount of competence with or mastery of relevant concepts. So Norman's knowledge of (1) counts as a priori and his knowledge of (2) does not.

I have some sympathy for this line of response. But I don't think it fully answers Williamson's argument. One issue is that the response relies on features that are specific to Williamson's example.[35] Consider the generic claim that ripe cranberries are red. It is plausible that becoming more skillful in visually identifying ripe cranberries contributes to mastery of the concept *ripe cranberry*. If that's so, then Williamson could simply replace "all recent volumes of Who's Who are red" with "ripe cranberries are red". Given this replacement, in both Norman's knowledge that all crimson things are red and his knowledge that ripe cranberries are red, experience contributes to mastery of the relevant concepts. That is so, even though Norman's knowledge that ripe cranberries are red is clearly a posteriori.

A second issue is that the response relies on a sharp distinction between what is and what is not relevant to the mastery of a concept. Being able to identify colors based on their appearance is supposed to count as relevant to conceptual mastery, while being able to identify recent volumes of *Who's Who* based on their appearance is not. How exactly should this distinction be drawn? The answer is not obvious. Without a clear answer, there is a worry that the distinction does not itself carve nature at its joints, and so presumably, the a priori/a posteriori distinction does not carve nature at its joints, either.

There is also a deeper issue with the proposed line of response. The distinction between what is and what is not part of conceptual mastery is a conceptual one, not an epistemic one. By contrast, the a priori/a posteriori distinction is an epistemic distinction. Why should a conceptual distinction matter for epistemology? It seems inappropriate to rely heavily on a conceptual distinction in drawing what should be an epistemic contrast, at least without an account of the

---

[34] A similar point could be made concerning our basic knowledge about the natural numbers. This knowledge relies on psychological mechanisms that are tied to perceptual experience – for instance, the mechanisms involved in object tracking and approximate cardinality comparisons. See Carey (2009).
[35] I owe this point to Elia Zardini (personal communication).

epistemic relevance of the conceptual distinction. And it is not immediately apparent why this distinction should be epistemically relevant.

What this last point suggests is that, if we are going to answer Williamson's argument by appealing to a disanalogy between Norman's knowledge of (1) and his knowledge of (2), the disanalogy we appeal to should be an epistemic one, or at least one of clear epistemic significance. In the next section, I will argue that there is such a disanalogy.

4. Responding to Williamson's Argument

Let's continue the two Norman cases. Suppose Norman comes to know, in the way Williamson sketches, that all crimson things are red. Suppose Norman then learns (misleading) information that strongly supports the claim that his training with the word 'crimson' – his being shown samples to which 'crimson' applies and does not apply and being told which was which – was a hallucination or a dream. Norman can still visually imagine a crimson object and identify the imagined object as red. Suppose after he learns this information, Norman retains his belief that all crimson things are red. Does this belief still count as knowledge?

For the most straightforward way of filling in the details of the case, the answer intuitively is yes. Norman still knows that all crimson things are red. He may lose his metalinguistic knowledge that the word 'crimson' applies to crimson things.[36] The information that discredited his training may defeat his knowledge of how the word is used in his linguistic community. So he may lose the knowledge that the word 'crimson' only applies to things that are red. But he still retains his knowledge that all crimson things are red. The process of visually imagining a crimson thing, identifying it as red, and (somehow) generalizing to all crimson things, would seem to suffice for that knowledge.[37]

By contrast, suppose Norman comes to know, in the way Williamson sketches, that all recent volumes of *Who's Who* are red. Suppose Norman then learns (misleading) information that strongly supports the claim that his training with the phrase 'recent volume of *Who's Who*' – his being shown samples to which 'recent volume of *Who's Who*' applies and does not apply and being told which was which – was a hallucination or a dream. Norman can still visually imagine a recent volume of *Who's Who* and identify the imagined volume as red. Suppose after he learns this information, Norman retains his belief that all recent volumes of *Who's Who* are red. Does this belief still count as knowledge?

For the most straightforward way of filling in the details of the case, the answer intuitively is no. The evidence Norman has that all recent volumes of *Who's Who* look a certain way has been defeated. He no longer has good reason to think that recent volumes of *Who's Who* are red, or

---

[36] I mean this claim de re and not de dicto.
[37] There are ways of filling out the case in which Norman also loses the knowledge that all crimson things are red. For instance, he might have background knowledge that he is hopeless in his thinking about colors without a great deal of training. But that's not the most natural way to develop the case.

even that any recent volume of *Who's Who* is red. So his belief that all recent volumes of *Who's Who* are red no longer counts as knowledge.

This disanalogy suggests that there is an important epistemic difference between the two cases.[38] This difference is not a difference in the psychological processes involved in Norman's coming to believe (1) and (2). Williamson is right that the way in which Norman came to believe (1) and the way in which he came to believe (2) are almost exactly alike. But that is compatible with there being a significant epistemic difference between the two cases. Williamson in effect assumes that any epistemically significant difference must stem from a difference in the cognitive processes underlying belief-formation. But this assumption is false. The fact that Norman's beliefs in (1) and (2) behave differently when Norman is provided information that discredits his earlier training strongly suggests that there is nonetheless a significant epistemic difference between the cases.

What is the epistemic difference between the cases? The natural suggestion to make is that the difference between the cases isn't a matter of a difference in the cognitive processes employed in belief-formation, but in the justification of the beliefs.[39] Consider the case of knowing some mathematical fact by coming to grasp a proof of the fact. In coming to grasp the proof, the thinker may rely on testimony, diagrams on paper, imagination, and the like. Such means involve experience – in reading or hearing the testimony, in constructing and viewing the diagram, and perhaps in tuning one's imagination. And this experience does not merely provide conceptual competence or contribute to conceptual mastery. Nevertheless, we should count the resulting knowledge as a priori. That is because the justification for the mathematical fact is the proof itself.[40] The fact that grasping the proof involves experience is neither here nor there.[41] The role experience plays in the process of belief-formation isn't directly relevant to whether a case of knowledge is a priori or a posteriori. What is directly relevant is the role, if any, that experience plays in the epistemic justification for the belief.[42]

What is the justification of Norman's belief in (1)? It is not completely clear what to say here. But there is a rough-and-ready diagnostic for when experience plays a role in the justification of a claim. Namely, if strong evidence that the experience is misleading defeats the justification of the claim, this provides (defeasible) reason to think that the experience is playing a genuinely

---

[38] Chalmers (2012, excursus eight), Jenkins and Kasaki (2015), and Melis and Wright (this volume) also suggest that there is an epistemic difference between the cases.

[39] Jenkins and Kasaki (2015) also make this suggestion.

[40] See Peacocke (2004, p. 25).

[41] As an aside, this point may help to explain a puzzling feature of Plato's *Meno*. In that dialogue, Socrates asks one of Meno's household slaves a series of questions that allows the slave to determine how to construct a square that has twice the area of a given square. Socrates argues that he did not teach the slave how to do this, but only asked questions, and so concludes that mathematical knowledge is innate. In the dialogue, Socrates asks what are obviously leading questions, so the argument strikes many as fallacious. But there is a response to this concern. Notice that if the slave were later to gain (misleading) evidence that Socrates is an unreliable testifier, this would not defeat his knowledge about how to construct a square that is twice the area of a given square. In that way, his mathematical knowledge behaves very differently from ordinary testimonial knowledge.

[42] A thinker's justification for a belief is not completely disconnected from the cognitive processes involved in the formation and retention of the belief (and the thinker's dispositions to give up the belief under certain circumstances). But it is a tricky question how exactly the two are related, a question that I cannot go into here.

justifying role.[43] If strong evidence that the experience is misleading does not defeat the justification of the claim, this provides (defeasible) reason to think that the experience is not playing a genuinely justifying role. In the case of Norman's belief that all recent volumes of *Who's Who* are red, evidence that discredits his experiences concerning recent volumes of *Who's Who* defeats his belief, and so there is reason to think that the justification of the belief relies on experience. In the case of Norman's belief that all crimson things are red, evidence that discredits his experiences of crimson things does not defeat his belief, and so there is reason to think that the justification of that belief does not rely on experience.

This is, I think, enough to show that Williamson's argument fails. But it would be good to have more to say about what exactly the justification is for Norman's beliefs in (1) and (2). Here is one potential view: The justification for Norman's belief that all recent volumes of *Who's Who* are red is (roughly) that (i) some (or better, those) recent volumes of *Who's Who* are red and (ii) recent volumes of *Who's Who* share the same color. Both claims are justified by experience, and so Norman's belief that all recent volumes of *Who's Who* are red is justified by experience. The justification for Norman's belief that all crimson things are red is very different. It is that crimson is a determinate of the determinable red. Experience does not play any role in justifying this claim, though it may play a role in Norman's coming to grasp this justification.

I don't want to rest much weight on any particular account of what the justification is for Norman's beliefs in (1) and (2). There are other options besides the one I have sketched.[44] The

---

[43] There are two reasons that this diagnostic only provides a defeasible reason. First, evidence that discredits an experience could also have some additional epistemic impact, and so might defeat the claim but not by way of discrediting the experience. Second, as Pryor (2013, sec. 4) points out, on many epistemological views, evidence can undermine one's justification for believing a claim by providing evidence that some other claim isn't true, without the initial justification including antecedent justification for that other claim. For instance, perhaps my justification for believing that there is a red apple in front of me can be undermined by evidence that my visual perception is not reliable without my initial justification for believing that there is a red apple in front of me having to include antecedent justification to believe that my visual perception is reliable. If that's right, then presumably there will be analogous cases in which evidence that discredits an experience serves to undermine justification for believing a claim without the initial justification for believing the claim depending on the experience.

[44] Melis and Wright (this volume, section 8) propose a different account of the crucial epistemic disanalogy between Norman's belief in (1) and his belief in (2). The disanalogy they point to concerns the move from specific claims – this [visually perceived] volume of *Who's Who* is red, and this [visually imagined] crimson thing is red – to general claims – all recent volumes of *Who's Who* are red, and all crimson things are red. On their view, for Norman to know (2) on the basis of knowing that this [visually perceived] volume of *Who's Who* is red, Norman must have antecedent warrant for the claim that recent volumes of *Who's Who* share the same color. And this warrant must be a posteriori. By contrast, for Norman to know (1) on the basis that this [visually imagined] crimson thing is red either (a) it is not required that Norman have antecedent warrant for the claim that the crimson object that Norman imagines is stereotypical of crimson objects, at least in terms of color, or (b) it is required that Norman have such antecedent warrant, but this warrant may be a priori. (Melis and Wright replace the crimson/red case with different case, but the substitution doesn't materially change things.)

Melis and Wright may be right about this disanalogy, but there is a worry about how far their view extends. Their diagnosis relies on the fact that (1) and (2) are general claims. But Williamson's two cases can be replaced with cases that do not obviously involve any move from specific to general claims. (This point was independently noticed by Elia Zardini (personal communication) and also appears in Williamson (this volume, footnote 2).) For instance, Williamson could instead use (1') this [visually imagined] crimson thing is red, and (2') this [visually perceived] volume of *Who's Who* is red. Alternatively, he could use (1'') some shade of red is a shade of crimson, and (2'') some red thing is a volume of *Who's Who*. And so on. Melis and Wright could try to extend their strategy to these cases by identifying a claim that Norman must have antecedent a posteriori warrant in to have knowledge of

more important point is that there is a plausible response to Williamson's argument against the naturalness and theoretical utility of the a priori/a posteriori distinction. Experience plays no role in the justification of Norman's belief that all crimson things are red. It may play a role in how Norman came to possess the ingredient concepts in that belief, in how he came to form the belief, or in how he came to grasp the justification for the belief, but it plays no role in the justification itself. So in answering Williamson's objection, the right distinction to draw is the distinction between cases of knowledge in which the justification of the belief involves experience, and cases of knowledge in which the justification of the belief does not.

5. Theoretical Utility

Assuming the discussion in the previous section is correct, Williamson's argument for the conclusion that the a priori/a posteriori distinction is not a natural one – that is, that it does not carve nature at its joints – is unsuccessful. But there remains the question of whether Williamson's conclusion is correct. Does the a priori/a posteriori distinction in fact carve nature at its joints? There also remains the question of whether the a priori/a posteriori distinction is a theoretically useful one. And these are both good questions.

As I mentioned above, the question of whether the a priori/a posteriori distinction is natural is distinct from the question of whether the distinction is a theoretically useful one. Naturalness is a metaphysical notion – whether a distinction is natural is a matter of how reality is structured. Theoretical utility is an epistemological or practical notion – whether a distinction is theoretically useful is a matter of whether drawing the distinction is useful for some epistemic purpose. So the two questions are at least notionally distinct. They are also, at least in principle, extensionally distinct. Presumably, any distinction that is natural is theoretically useful – it is epistemically valuable to recognize a natural distinction and to investigate what falls on either side of it.[45] But the other direction seems to fail. There can be distinctions that do not carve nature at its joints, but which are theoretically useful or useful in some other way.[46] Consider, for example, the concept *salad*. Presumably, the salad/non-salad distinction does not carve at the (culinary?) joints. There is little in common between a green salad, a composed salad, a three-bean salad, tuna salad, pasta salad, fruit salad, and so forth.[47] There are also examples of salads that are very similar to non-salads. So there is good reason to think that the salad/non-salad distinction is not a natural one. However, the concept of a salad does have practical utility –

(2') (or (2'') or …), while arguing that either (a) Norman need not have antecedent warrant in the corresponding claim to have knowledge of (1') (or (1'') or …), or (b) Norman does need have such antecedent warrant to have knowledge of (1') (or (1'') or …) , but this warrant may be a priori. The trouble is that I don't see how to implement this strategy for those cases.

[45] See Sider (2011, ch. 4.5) for defense of a strong version of this claim.

[46] Here I disagree with Hawthorne, who writes "If an epistemological distinction fails to carve at the epistemological joints, then it is not worthy of serious and protracted discussion." Hawthorne (2007, p. 201)

[47] Reliable informants have told me that American and British English differs in whether some of these count as kinds of salad, which is itself somewhat telling.

when ordering a meal at a restaurant, sometimes what one wants is a salad. So the salad/non-salad distinction is a practically useful one, despite not being natural.[48]

How can we determine if the a priori/a posteriori distinction is theoretically useful? To get purchase on this question, it helps to ask a more basic question: What is the purpose (or function or point) of drawing the a priori/a posteriori distinction? That is, what is the distinction for?[49]

It is important to recognize that the terms "a priori" and "a posteriori" are technical terms, introduced (or, rather, repurposed) for a particular theoretical aim. As Daniele Sgaravatti writes, "[I]t should not be forgotten that 'a priori' and 'a posteriori' are not ordinary terms in English or Latin, and there are no ordinary terms in any other language that express the same concepts. They are technical terms—pieces of philosophical jargon, or theoretical tools."[50]

What is the purpose, then, of drawing the a priori/a posteriori distinction? There is a plausible suggestion to make here, namely, that the purpose of drawing the distinction is to identify epistemological problem cases. Indeed, this is the use to which Kant puts the distinction. In Kant's view, we understand (more-or-less) how knowledge based on experience works. What is more mysterious is how we can have genuine knowledge not based on experience, such as knowledge of mathematics or of metaphysical claims. The role of the concept of the a priori is to identify puzzle cases so we can focus philosophical attention on them. This is (close to) explicit in Kant's writing, and is a perennial theme in the literature on the a priori.[51]

So the proposal, then, is this: We understand – more-or-less – the epistemology of simple perceptual knowledge. The epistemology of non-perceptual knowledge is far less clear. The purpose of labeling a case of knowledge as a priori is to claim that its epistemology should not be assimilated to the epistemology of perception. Instead, it is something of a puzzle case.[52]

---

[48] For discussion of the salad example, see Dorr and Hawthorne (2014). Food-based examples are surprisingly good for illustrating this point. The concept *soup* and the concept *sandwich* are also non-joint-carving but useful. Indeed, there is an online computer game, "Something Something Soup Something", whose point seems in part to be that the concept *soup* is not a natural one.

[49] The idea of making epistemological progress by asking questions about the function of our epistemic notions is due, in recent philosophy, to Craig (1990), which raises the question of the purpose of knowledge attributions.

[50] Sgaravatti (2020, p. 1134).

[51] There are alternative suggestions about the point of drawing the a priori/a posteriori distinction that could be made. The only other one that strikes me as worth considering is that the point of drawing the distinction is to help us to identify parts of inquiry where no further evidence needs to be gathered. On this proposal, the idea is that the distinction has a practical purpose. If we find out some matter is a priori, we won't waste our time looking for empirical evidence about the matter. I do think that there is a useful notion in the ballpark here. But, first, this is not how the term "a priori" is typically deployed. I can't think of a case in which an author argues that some matter is a priori for the purpose of arguing that we needn't gather more empirical evidence about it. And, second, the most useful notion in the ballpark is the notion of possessing enough evidence on some matter so that it wouldn't be useful to gather any more. There are a posteriori claims about which it wouldn't be useful to spend any more time gathering evidence. And there are a priori claims about which it would be useful to gather more evidence. So, the right distinction to play this role doesn't closely resemble the a priori/a posteriori distinction.

[52] Of course, there are puzzles concerning perceptual knowledge, especially concerning how best to answer skeptical challenges. But the a priori is more puzzling than perceptual knowledge. Indeed, a central motivation for Empiricism as opposed to Rationalism is that the claim that we can have priori knowledge of the structure of reality is too puzzling to accept.

There are two complications that should be mentioned here. The first is that Kant thought that some cases of a priori knowledge were unproblematic – namely, knowledge of analytic truths. For him, the real puzzle concerned synthetic a priori knowledge. But this complication doesn't tell against the proposal. For Kant, the purpose of drawing the a priori/a posteriori distinction is to narrow down cases of knowledge to a superset of the puzzle cases. We can then use the analytic/synthetic distinction to narrow down the cases still further. So on Kant's view, the point of the distinction really is to help us to identify puzzle cases. Moreover, Kant was too quick to assume that our knowledge of analytic truths is unproblematic. Famously, W. V. Quine's "Two Dogmas of Empiricism" provided a sustained attack on the analytic/synthetic distinction.[53] But even putting aside Quine's objections, it is mysterious how we can acquire knowledge of a claim just because the claim is (somehow) built into the meaning of our words or the nature of our concepts.[54] So the entire category of a priori knowledge is puzzling.

The second complication is that non-a priori knowledge is not automatically unproblematic. The definition of "a posteriori" permits a posteriori knowledge to have an a priori component. A priori knowledge cannot depend on a posteriori knowledge, but a posteriori knowledge can depend on a priori knowledge. Indeed, it is plausible that in many (most? all?) cases, it will. For instance, if having inductive knowledge requires antecedently knowing (or having justification to believe) that the world contains projectible regularities, and if this knowledge (or justification) is a priori, then inductive knowledge will always depend on some a priori knowledge (or justification). So if a priori knowledge is puzzling, inductive knowledge will be puzzling, too.

This complication also doesn't tell against the proposal. In identifying puzzle cases to focus one's attention on, it can be helpful to identify the "cleanest" cases of the puzzling phenomenon – cases that are the best exemplars of the puzzling phenomenon and that have the fewest distracting features. Cases of a priori knowledge are like that. They are puzzle cases that avoid the distracting features that a tie to experience can bring. Part of the point of identifying cases of a priori knowledge is to help us to identify the best cases to consider for making philosophical progress.

Supposing the proposal about the purpose of drawing the a priori/a posteriori distinction is correct, then (something like) the category of the a priori is theoretically useful after all, at least for epistemologists. It is useful to identify cases of knowledge in which it is especially puzzling what the explanation of our knowledge is. Identifying such cases enables us to focus attention on pressing epistemological tasks, and (one hopes) to make progress in epistemology.

6. Updating the Distinction

My proposal about the purpose of the a priori/a posteriori distinction has an important implication. There are multiple ways in which a case of knowledge may be disanalogous to simple cases of perceptual knowledge. There are two disanalogies that are perhaps the most

---

[53] Quine (1951).
[54] See Boghossian (1997) for relevant discussion.

important: the explanation of how the belief is justified may be different, and the explanation of how the thinker came to have a true belief may be different.

Consider a simple case of perceptual knowledge – say, my knowledge that there is a red apple in front of me. It is plausible that in this case, the explanation of how my belief is justified centrally involves phenomenality – I am justified in believing that there is a red apple in front of me in part because it visually appears to me that there is a red apple in front of me.[55] It is also plausible that in this case, the explanation of how it is that I have a true belief centrally involves a causal relation – I have a true belief that there is a red apple in front of me in part because the presence of the red apple caused me to believe that there is a red apple in front of me. And, presumably, the same is true for other simple cases of perceptual knowledge.

If this is right, then there are (at least) two important ways in which a case of knowledge can be disanalogous to a simple case of perceptual knowledge: (i) the justification of the belief does not involve phenomenality, and (ii) the belief does not stand in a causal relation to what the belief is about. These disanalogies raise distinct epistemological puzzles: How can our beliefs about some subject matter be justified if not ultimately (at least in part) in terms of how things phenomenally seem to us?[56] How is it that we get things right about the subject matter if our beliefs do not stand in some causal relation to what they are about?

So there are at least two different kinds of epistemological puzzles to solve in the ballpark of the a priori/a posteriori distinction. If the purpose of drawing the distinction is to help us to identify puzzle cases, what this suggests is that we should draw (at least) two distinctions – one concerning phenomenality and one concerning causation.[57]

This way of thinking about the a priori and a posteriori has several nice features. First, when we consider domains that are typically taken to be a priori – for instance, logic, mathematics, fundamental morality, and the like – there are two important kinds of epistemological questions that arise. (i) How is it that our beliefs about the domain are justified? And (ii) how is it that we have true beliefs about the domain? Both questions are important. (The latter question is the focus of the Benacerraf-Field objection in the philosophy of mathematics and evolutionary

---

[55] This is not sufficient for me to be justified: I must presumably lack defeaters for the belief, and perhaps must also have justification to believe that how the world appears is an at least somewhat reliable guide to how the world in fact is. See Pryor (2000) and Huemer (2007) for the view that its perceptually seeming to a thinker that p yields pro tanto justification for the thinker to believe that p. See Wright (2002) and White (2006) for the view that the thinker must also antecedently have warrant for (something like) the claim that how things perceptually seem is a reliable guide to how things are. A phenomenal-based account of perceptual justification is not the only kind of account that is available. But it is one I am sympathetic to, and in what follows, I will assume that it is correct. (Many of the claims below would still obtain given an alternative account of perceptual justification, mutatis mutandis.)

[56] This puzzle is particularly pressing if we have a conception of justification that ties it to epistemic responsibility. How can it be responsible to have beliefs that are not based in how things seem to us?

[57] Someone might argue that the distinction concerning phenomenality is the genuine a priori/a posteriori distinction and that the distinction concerning causation is not, or vice versa. I have no real quarrel with either way of describing things. My main point is that both distinctions are theoretically useful. How exactly we should label them is of lesser importance.

debunking arguments in metaethics.)[58] So splitting the traditional a priori/a posteriori distinction into two distinctions – one focused on justification and the other focused on truth – fits well with the epistemological terrain.

Second, recall the question of how we should understand "experience" in the definition of a priori knowledge as independent of experience. There were several options, but all of them faced difficulties. The present suggestion is that we don't need to choose. There are (at least) two ways to understand what experience is: (i) as mental states that have a phenomenal aspect; and (ii) as mental states that are caused by what they are about.[59] There is an a priori/a posteriori distinction for each of them. On both distinctions, knowledge falls on the a priori side if the knowledge is independent of experience, but in one case "experience" is understood to involve phenomenality and in the other case "experience" is understood to involve causation. (As we'll see below, there is reason to liberalize the second distinction a bit further.)

Third, this way of thinking about the a priori/a posteriori distinction can help us to better categorize some confusing cases. Consider a case of true innate belief.[60] For instance, consider our belief that objects tend to have continuous trajectories in space.[61] Such a belief is plausibly part of the innate cognitive system involved in object representation that we evolved to have via natural selection. It is plausible that such a belief does – or could – count as genuine knowledge.[62] If so, is such knowledge a priori or a posteriori? On the one hand, the belief does not seem to depend on phenomenal experience. On the other hand, it seems problematic to count knowledge about the typical behavior of physical objects as a priori.

The proposal that we split the a priori/a posteriori distinction into two distinctions helps us to handle this case. The belief is a priori in the sense that its justification does not depend on phenomenality. However, it is a posteriori in the sense that the explanation of how we ended up with a true belief does involve a causal relationship between the belief and what the belief is about (mediated by evolution by natural selection). Moreover, this categorization helps us to identify the particularly puzzling aspect of such a belief – if the explanation of how the belief is justified does not involve phenomenality, just what is the explanation of its justification?

Similarly, on some views of the justification of beliefs about mathematics, morality, and metaphysics (etc.), such beliefs are justified by rational insights, which are understood to be phenomenally-laden cognitive states.[63] Supposing we can have knowledge in this way, should such knowledge count as a priori or a posteriori? On the one hand, logical, mathematical, and ethical knowledge seem to be paradigm cases of a priori knowledge. On the other hand, if the proposed account is correct, it is difficult to see why beliefs justified by phenomenal states

---

[58] See the Introduction and title essay of Field (1989) for the canonical presentation of the Benacerraf-Field argument. See Street (2006) and Joyce (2006) for evolutionary debunking arguments. See Schechter (2018) for a discussion of both kinds of arguments, focused on the case of morality.

[59] To be plausible, both suggestions need significant refinement.

[60] This kind of case plays an important role in Hawthorne (2007).

[61] See Spelke and Kinzler (2007).

[62] For discussion of innate a priori knowledge, see Cutter (2020) and Mackie (1969/1970).

[63] See, for instance, Bengson (2015) and Chudnoff (2011). For the record, I am suspicious of this kind of account.

shouldn't be assimilated to cases of perceptual justification. Indeed, rational insight-based accounts are often motivated by drawing a close analogy with perceptual justification.

The proposal that we split the a priori/a posteriori distinction into two distinctions helps us to handle these cases, too. Such beliefs are a posteriori in the sense that their justification depends on phenomenality. However, they are a priori in the sense that the explanation of how we ended up with a true belief does not involve a causal relationship between the belief and what the belief is about. Moreover, this categorization helps us to identify the particularly puzzling aspect of such beliefs, on a rational insight-based view – if the explanation of how we ended up with true beliefs does not involve a causal relationship between the belief and what the belief is about, just what is the explanation of how we ended up with true beliefs?

If all of this is correct, then there is a long-standing mistake (or, perhaps better, omission) in the literature. It is a mistake to think about the a priori/a posteriori distinction as though there is only one distinction to draw. Since the purpose of drawing the distinction is to help us to identify puzzle cases, and there are different kinds of puzzles, it is theoretically useful to draw multiple distinctions, one per kind of puzzle. Otherwise, there is the danger that we will conflate distinct kinds of puzzles, and that we will provide an account of how a priori knowledge is possible about some domain that solves the wrong puzzle. For instance, we may be tempted to appeal to rational insight to explain our knowledge of some domain, when the central epistemological problem concerning that domain is to explain how we get it right and not how we are justified. We may be tempted to appeal to evolution by natural selection to explain our knowledge of some other domain, when the central epistemological problem concerning that domain is to explain how we are justified and not how we get it right.

This is an instance of a more general point about philosophical methodology. The ways of thinking we inherited from our intellectual forebears should always be open to challenge. It can be useful to consider the purposes of the distinctions we inherited from them and whether there might be a more effective way to fulfill those purposes.

Before I conclude, let me say just a little bit more about the two distinctions, and (tentatively) propose terminology for them. Let me start with the distinction concerning phenomenality. The idea is that we should distinguish between those cases of knowledge in which the explanation of the justification of the belief involves phenomenality and those cases of knowledge in which it does not. Really, though, the distinction needn't be restricted to cases of knowledge. We can distinguish between cases of justified beliefs (or rational credences) in which the explanation of the justification of the belief (or rationality of the credence) depends in part on phenomenality and cases in which the explanation does not depend on phenomenality. We can call this the distinction between "phenomenally dependent" and "phenomenally independent" justified beliefs.[64]

---

[64] We can generalize this distinction still further. A belief may have two independent justifications, one involving phenomenality and one not. So it would be better to draw the distinction not between beliefs but between justifications for beliefs. We can also ask about the justification of things other than beliefs, such as the justification for employing rules of inference and belief-forming methods more generally. So it would be better still to draw the distinction between justifications.

The second distinction concerns causation. The idea is that we should distinguish between those cases of knowledge in which the explanation of why the thinker came to have a true belief involves the claim that the belief stands in a causal relation to what the belief is about and those cases of knowledge in which it does not. (Recall the case of the red apple: my belief that there is a red apple was caused by the presence of the red apple.) Really, though, the distinction doesn't need to concern cases of knowledge. It is more general to distinguish between cases of true beliefs (or accurate credences) in which the explanation of why the thinker has a true belief depends in part on a causal relation between the belief and what the belief is about and cases in which the explanation does not. We can call this the distinction between "causally dependent" and "causally independent" true beliefs.[65]

There are two clarifications that should be made about the causally dependent/causally independent distinction. First, the causal relationship between the belief and what the belief is about needn't be that the belief is caused by what it's about. The belief the p and the fact that p may share a common cause. Alternatively, the belief that p may itself cause the fact that p. Such cases should go on the causally dependent side of the distinction. That's because the point of making the distinction is to help to identify puzzle cases – cases in which it is at least somewhat mysterious how the thinker has a true belief. There is no mystery when the belief that p and the fact that p have a common cause, or when the belief that p causes the fact that p. Second, the causal relationship between the belief that p and the fact that p needn't be immediate, or even be one that took place solely during the thinker's lifespan. Recall the case of the innate belief that physical objects tend to have continuous trajectories. If the explanation of how we have a true belief in this claim involves selective pressures on our distant ancestors, the fact that physical objects tend to have continuous trajectories in part caused our belief that they do, and the explanation of our reliability is unmysterious. So it again makes sense to locate this case on the causally dependent side of the distinction.

To summarize:

> A justified belief is *phenomenally dependent* just in case the explanation of the justification of the belief depends in part on phenomenality. A justified belief is *phenomenally independent* just in case the explanation of the justification of the belief does not depend in part on phenomenality.

> A true belief is *causally dependent* just in case the explanation of why the thinker has a true belief depends in part on a causal relation between the belief and what the belief is about. A true belief is *causally independent* just in case the explanation of why the

---

[65] We can generalize this distinction still further. Even if a belief is false, it may have verisimilitude. That is, it may be close to true. (Consider a belief in Newton's laws.) We could distinguish between cases in which the explanation of how a thinker came to have a belief with verisimilitude involved a causal relation with the topic of the belief, and a case in which it did not. We can also ask about the truth-conduciveness of things other than beliefs, such as the truth-conduciveness of a rule of inference or belief-forming method.

thinker has a true belief does not depend in part on a causal relation between the belief and what the belief is about.[66]

These two distinctions can, at least in principle, cut across one another. So a justified true belief can be in any of four categories: (i) both phenomenally and causally independent, (ii) phenomenally dependent and causally independent, (iii) phenomenally independent and causally dependent, and (iv) both phenomenally and causally dependent.

There are potential examples of justified true belief in all four categories. Take a simple logical, mathematical, or conceptual belief – for instance, everything is self-identical, 2 is greater than 1, or bachelors are unmarried. Such beliefs are plausibly both phenomenally and causally independent. That is why such beliefs are so puzzling.

As I mentioned above, on some views, at least some logical, mathematical, conceptual, and philosophical beliefs are justified on the basis of phenomenally-laden cognitive states, called rational insights. Such beliefs, when justified and true, are presumably phenomenally dependent and causally independent. They are justified via phenomenality, but they are not caused by what they are about. (How could a mathematical belief be caused by the relevant mathematical fact, at least assuming mathematical Platonism?). If there are beliefs in this category, what is most puzzling about them is not how they are justified but the explanation of how the relevant thinker ended up with true beliefs.

The category of phenomenally independent and causally dependent beliefs plausibly includes justified true innate beliefs such as the belief that physical objects typically follow continuous paths. What is puzzling about such beliefs is how they are justified, since phenomenality doesn't play any obvious part in their justification.

Finally, ordinary empirical beliefs, such as my beliefs that it is currently raining, that the word "table" stands for a kind of furniture, and that general relativity well describes the behavior of highly massive objects, fall into the category of beliefs that are both phenomenally and causally dependent. There are puzzles about how such beliefs are justified, and how we ended up with true beliefs. But such puzzles are not nearly as pressing as the epistemological puzzles concerning the other three categories.

There is one final issue remaining. I've argued, contra Williamson and Hawthorne, that the a priori/a posteriori distinction, when properly understood – and split into two distinctions – is theoretically useful. But I have not discussed the question of whether these distinctions carve at the joints. Are they natural distinctions? What should we think about that question?

---

[66] Notice that for each of these two distinctions, there is an analogue of the distinction between merely enabling and genuinely evidential roles. For a justified belief to be phenomenally dependent, phenomenality must contribute to the justification of the belief, not merely enable the belief to be justified. For a true belief to be causally dependent, what the belief is about must be a cause of the belief (or an effect of the belief, or an effect of a common cause, etc.), not merely a causal enabling condition. I owe this observation to Eileen Nutting.

Naturalness comes in degrees. The central question for us here is not whether the two distinctions are maximally (i.e., "perfectly") natural, but whether they have a high degree of naturalness, that is, whether they are reasonably natural.

The distinction between the Ps and the Qs is a (reasonably) natural distinction just in case either the property of being P or the property of being Q, or both, is a (reasonably) natural property. For a property to be (reasonably) natural requires that things that possess the property be (reasonably) similar to one another, and that they be (reasonably) dissimilar from things that do not possess the property. There should be some sort of deep commonality between things that possess the property and a deep difference between things that do and things that don't possess the property.[67]

For instance, the distinction between electrons and non-electrons is a natural distinction because the property of being an electron is a natural property. The property of being an electron is a natural property because electrons closely resemble one another, and there is an important difference between electrons and non-electrons.[68]

Are the phenomenally dependent/phenomenally independent and the causally dependent/causally independent distinctions natural ones? I think at our present state of understanding we should be agnostic. The purpose of the distinctions is to identify cases to focus attention on to make philosophical progress. This is compatible with the distinctions being natural and compatible with the distinctions being non-natural. What would have to be determined is whether the properties of being phenomenally dependent, phenomenally independent, causally dependent, and causally independent are natural properties. Do the classes of beliefs having these properties share a deep commonality? Is there is a deep difference between beliefs that do and beliefs that don't have one of these properties? I don't think we yet know the answers to these questions. We don't yet know, for instance, if the phenomenally independent beliefs are highly heterogenous or not. So to answer the question about the naturalness of the distinctions, we must make further progress in epistemology.[69]

References

Adams, Robert Merrihew (1994). *Leibniz: Determinist, Theist, Idealist*. New York: Oxford University Press.

Aristotle (1994) [4th century BCE]. *Posterior Analytics*, translated and with a commentary by Jonathan Barnes, second edition. Oxford: Oxford University Press.

---

[67] See Dorr (2019) for discussion of natural properties.

[68] Notice that the property of being a non-electron is not a natural property – the non-electrons are a motley – but this is compatible with the distinction between electrons and non-electrons being a natural distinction.

Arnauld, Antoine and Nicole, Pierre (1996) [1662/1683]. *Logic or the Art of Thinking*, translated by Jill Vance Buroker. Cambridge: Cambridge University Press.

Baumgarten, Alexander Gottlieb (2013) [1739]. *Metaphysics*, translated and edited by Courtney Fugate and John Hymers. Bloomsbury. London: Bloomsbury.

Baumgarten, Alexander Gottlieb (1761). *Acroasis Logica*. Halle.

Bengson, John (2015). "The Intellectual Given", *Mind* 124(495): 707–760.

Block, Ned (1995). "On a Confusion about a Function of Consciousness", *Brain and Behavioral Sciences* 18(2): 227–247.

Boghossian, Paul (1997). "Analyticity", in Bob Hale and Crispin Wright (eds.), *A Companion to the Philosophy of Language*. Oxford: Blackwell. pp. 331–368.

Boghossian, Paul (2020). "Do We Have Reason to Doubt the Importance of the Distinction between A Priori and A Posteriori Knowledge?", in Paul Boghossian and Timothy Williamson, *Debating the A Priori*. Oxford: Oxford University Press. pp. 137–155.

BonJour, Laurence (1998). *In Defense of Pure Reason*. Cambridge: Cambridge University Press.

Burge, Tyler (1993). "Content Preservation", *Philosophical Review* 102(4): 457–488.

Carey, Susan (2009). *The Origin of Concepts*. New York: Oxford University Press.

Casullo, Albert (2003). *A Priori Justification*. New York: Oxford University Press.

Casullo, Albert (2011). "Reply to My Critics: Anthony Brueckner and Robin Jeshion", in Michael J. Shaffer and Michael Veber (eds.), *What Place for the A Priori*? Chicago: Open Court, pp. 111–136.

Casullo, Albert (this volume). "A Defense of the Significance of the A Priori – A Posteriori Distinction".

Chalmers, David (2012). *Constructing the World*. Oxford: Oxford University Press.

Chudnoff, Elijah (2011). "The Nature of Intuitive Justification", *Philosophical Studies* 153(2): 313–333.

Coffa, Alberto (1991). *The Semantic Tradition from Kant to Carnap: To the Vienna Station*. Cambridge: Cambridge University Press.

Cohen, Stewart (2010). "Bootstrapping, Defeasible Reasoning, and A Priori Justification", *Philosophical Perspectives* 24, *Epistemology*, 141–159.

Craig, Edward (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Oxford University Press.

Cutter, Brian (2020). "Why Nearly Everything Knowable is A Priori", *Philosophy and Phenomenological Research* 101(1): 80–100.

Dorr, Cian (2019). "Natural Properties", *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2019/entries/natural-properties/>.

Dorr, Cian and Hawthorne, John (2014). "Semantic Plasticity and Speech Reports", *The Philosophical Review* 123(3): 281-338.

Evans, Gareth (1979). "Reference and Contingency", *The Monist* 62(2): 161–189.

Field, Hartry (1989). *Realism, Mathematics and Modality*. Oxford: Blackwell.

Frege, Gottlob (1974) [1884]. *The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number*, translated by J.L. Austin. Oxford: Blackwell, second revised edition.

Hawthorne, John (2007). "A Priority and Externalism", in Sanford C. Goldberg (ed.), *Internalism and Externalism in Semantics and Epistemology*. Oxford: Oxford University Press. pp. 201–218.

Huemer, Michael (2007). "Compassionate Phenomenal Conservatism", *Philosophy and Phenomenological Research* 74(1): 30–55.

Hume, David, (2007) [1748]. *An Enquiry concerning Human Understanding*, Peter Millican (ed.). Oxford: Oxford University Press.

Jenkins, C. S. I. and Kasaki, Masashi (2015). "The Traditional Conception of the A Priori", *Synthese* 192(9): 2725–2746.

Jeshion, Robin (2011). "Experience as a Natural Kind: Reflections on Albert Casullo's A Priori Justification", in Michael J. Shaffer and Michael Veber (eds.), *What Place for the A Priori*? Chicago: Open Court, pp. 93–109.

Joyce, Richard (2006). *The Evolution of Morality*. Cambridge, Mass.: MIT Press.

Kant, Immanuel (1998) [1781/1787]. *Critique of Pure Reason*, translated by Paul Guyer and Allen Wood. Cambridge: Cambridge University Press.

Kripke, Saul (1980). *Naming and Necessity*. Cambridge, Mass.: Harvard University Press.

Lambert, Johann Heinrich (1764). *Neues Organon*. Leipzig.

Lewis, David (1980). "Mad Pain and Martian Pain", in Ned Block (ed.), *Readings in the Philosophy of Psychology*, Volume I. Cambridge, Mass.: Harvard University Press, pp. 216–222.

Lewis, David (1983). "New Work for a Theory of Universals", *Australasian Journal of Philosophy* 61(4): 343–377.

Mackie, J. L. (1969/1970). "The Possibility of Innate Knowledge", *Proceedings of the Aristotelian Society* 70: 245–257.

McGinn, Colin (1975/1976). "A Priori and A Posteriori Knowledge", *Proceedings of the Aristotelian Society* 76: 195–208.

Melis, Giacomo and Wright, Crispin (this volume). "Williamsonian Scepticism about the A Priori".

Mittelstrass, Jürgen (1977). "Changing Concepts of the A Priori", in Robert Butts and Jaakko Hintikka (eds.), *Historical and Philosophical Dimensions of Logic, Methodology and Philosophy of Science*. Dordrecht: Springer, pp. 113–128.

Nagel, Thomas (1974). "What is it Like to be a Bat?", *Philosophical Review* 83(4): 435–450.

William of Ockham (1974) [c. 1323]. *Opera Philosophica I: Summa Logicae*, edited by Philotheus Boehner, Gedeon Gál, and Stephen Brown. Bonaventure, New York: Franciscan Institute.

Peacocke, Christopher (2004). *The Realm of Reason*. Oxford: Oxford University Press.

Plato (2011) [4th century BCE]. *Meno and Phaedo*, David Sedley (ed.), and Alex Long (ed., tr.). Cambridge: Cambridge University Press.

Plato (2012) [4th century BCE]. *Republic*, translated by Christopher Rowe. London: Penguin Classics.

Pryor, James (2000). "The Skeptic and the Dogmatist", *Noûs* 34(4): 517–549.

Pryor, James (2013). "Problems for Credulism", in Chris Tucker (ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. New York: Oxford University Press, pp. 89–131.

Quine, Willard V. O. (1951). "Two Dogmas of Empiricism", *Philosophical Review* 60(1): 20–43.

Quine, Willard V. O. (1969). "Natural Kinds", in Nicholas Rescher (ed.), *Essays in Honor of Carl G. Hempel*. Dordrecht: Springer, pp. 5–23.

Schechter, Joshua (2017). "No Need for Excuses: Against Knowledge-First Epistemology and the Knowledge Norm of Assertion", in J. Adam Carter, Emma Gordon, and Benjamin Jarvis (eds.), *Knowledge-First: Approaches in Epistemology and Mind*. Oxford: Oxford University Press. pp. 132–159.

Schechter, Joshua (2018). "Explanatory Challenges in Metaethics", in Tristram McPherson and David Plunkett (eds.), *Routledge Handbook of Metaethics*. New York: Routledge. pp. 443–459.

Schepers, Heinrich, Tonelli, Giorgio and Eisler, Rudolf (1971). "A Priori/A Posteriori", in Joachim Ritter (ed.), *Historisches Wörterbuch der Philosophie*, Part I: A-C. Basel: Schwabe.

Sgaravatti, Daniele (2020). "Experience and Reasoning: Challenging the A Priori/A Posteriori Distinction", *Synthese* 197(3): 1127–1148.

Sider, Theodore (2011). *Writing the Book of the World*. Oxford: Oxford University Press.

Silins, Nicholas (2008). "Basic Justification and the Moorean Response to the Skeptic", *Oxford Studies in Epistemology* 2: 108–140.

Smit, Houston (2009). "Kant on Apriority and the Spontaneity of Cognition", in Samuel Newlands and Larry M. Jorgensen (eds.), *Metaphysics and the Good: Themes from the Philosophy of Robert Merrihew Adams*. Oxford: Oxford University Press, pp. 188–251.

Spelke, Elizabeth and Kinzler, Katherine (2007). "Core Knowledge", *Developmental Science* 10(1): 89–96.

Street, Sharon (2006). "A Darwinian Dilemma for Realist Theories of Value", *Philosophical Studies* 127(1): 109–166.

Wedgwood, Ralph (2013). "A Priori Bootstrapping", in in Albert Casullo and Joshua C. Thurow (eds.), *The A Priori in Philosophy*. Oxford: Oxford University Press, pp. 226–246.

White, Roger (2006). "Problems for Dogmatism", *Philosophical Studies* 131(3): 525–557.

Williamson, Timothy (2007). *The Philosophy of Philosophy*. Oxford: Wiley-Blackwell.

Williamson, Timothy (2013). "How Deep is the Distinction between A Priori and A Posteriori Knowledge?", in Albert Casullo and Joshua C. Thurow (eds.), *The A Priori in Philosophy*. Oxford: Oxford University Press, pp. 291–312.

Williamson, Timothy (this volume). "More Williamsonian Scepticism about the A Priori – A Posteriori Distinction".

Wolff, Christian (1728). *Philosophia Rationalis sive Logica.* Frankfurt.

Wright, Crispin (2002). "(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell", *Philosophy and Phenomenological Research* 65(2): 330–348.