

How Deep is the Distinction between A Priori and A Posteriori Knowledge?¹

Timothy Williamson

Abstract: The paper argues that, although a distinction between a priori and a posteriori knowledge (or justification) can be drawn, it is a superficial one, of little theoretical significance. The point is not that the distinction has borderline cases, for virtually all useful distinctions have such cases. Rather, it is argued by means of an example, the differences even between a clear case of a priori knowledge and a clear case of a posteriori knowledge may be superficial ones. In both cases, experience plays a role that is more than purely enabling but less than strictly evidential. It is also argued that the cases at issue are not special, but typical of a wide range of others, including knowledge of axioms of set theory and of elementary logical truths. Attempts by Quine and others to make all knowledge a posteriori ('empirical') are repudiated. The paper ends with a call for a new framework to be developed for analysing the epistemology of cognitive uses of the imagination.

Keywords: apriori, aposteriori, knowledge, imagination, logic, mathematics

1. The distinction between a priori and a posteriori knowledge can be introduced the bottom-up way, by examples. I know a posteriori whether it is sunny. I know a priori that if it is sunny then it is sunny. Such examples are projectible. We learn from them how to go on in the same way, achieving fair levels of agreement in classifying new cases without collusion. Of course, as well as clear cases on each side there are unclear cases, which elicit uncertainty or disagreement when we try to classify them as a priori or as a posteriori. But virtually all useful distinctions are like that. If we want to be more precise, we can stipulate a sharper boundary with the clear cases of the a priori on one side and the clear cases of the a posteriori on the other. How could such a distinction be problematic? If some philosopher's theory puts a clear case on the wrong side of the line, surely that is a problem for the theory, not for the distinction.

The risk for the bottom-up method of introduction is that it may make a distinction of no special significance. On that scenario, our classifications follow similarities and differences that, although genuine, are largely superficial, like a taxonomy of plants and animals based only on colour. If so, epistemologists would do better to avoid the distinction between the a priori and the a posteriori in their theorizing, because it distracts them from deeper similarities and differences.

The alternative method of introduction is top-down, by a direct statement of the difference between the a priori and the a posteriori in epistemologically significant theoretical terms. For instance: a priori knowledge is independent of experience; a posteriori knowledge depends on experience. The risk for the top-down method is that it may turn out that everything is on the same side of the theoretically drawn line. If Quine (1951) is right, all knowledge depends on experience. As with the risk for the bottom-up

method, that would make the distinction epistemologically useless, but for a different reason.

Friends of the distinction typically assume that the bottom-up and top-down methods yield equivalent results and so are mutually supporting, each averting the risk to the other. Of course, the risk here is that the two methods may have incompatible results. If so, assuming otherwise leads us into epistemological error.

But are any of the risks realized? In this paper I will suggest two ways in which reliance on a distinction between the a priori and the a posteriori does more harm than good in epistemology. Neither of them is exactly Quine's. As a framework for discussion, I first sketch the top-down distinction as contemporary philosophers tend to conceive it.

2. The distinction between the a priori and the a posteriori is primarily a classification of specific ways of knowing.² A way of knowing is a priori if and only if it is independent of experience. It is a posteriori if and only if it depends on experience. The relevant senses of 'independent' and 'experience' are discussed below. Every specific way of knowing is either a priori or a posteriori, and not both. One knows p a priori if and only if one knows p in an a priori way. One knows p a posteriori if and only if one knows p in an a posteriori way. Thus if one knows p , one knows it either a priori or a posteriori.

One may know p both a priori and a posteriori, if one knows it in several ways, some a priori, some a posteriori. Tradition excluded that case on the grounds that only necessities (truths that could not have been otherwise) are known a priori whereas only contingencies (truths that could have been otherwise) are known a posteriori. But that was a mistake. Here is a simple counterexample. Suppose that Mary is good at

mathematics but bad at geography, while John is bad at mathematics but good at geography. Both of them can perform elementary deductions. Mary knows a priori by the usual standards that $289 + 365 = 654$ and does not know at all that there are cable cars in Switzerland. John knows a posteriori by the usual standards that there are cable cars in Switzerland but does not know at all that $289 + 365 = 654$. From the premise that $289 + 365 = 654$, Mary competently deduces the disjunctive conclusion that either $289 + 365 = 654$ or there are cable cars in Switzerland (since a disjunction follows from either disjunct), and thereby comes to know the disjunction a priori by the usual standards, since the logical deduction introduces no dependence on experience. Meanwhile, from the premise that there are cable cars in Switzerland, John competently deduces the same disjunctive conclusion, and thereby comes to know it a posteriori by the usual standards, for although the disjunction itself is a priori, his knowledge of the conclusion inherits the dependence on experience of his knowledge of the premise. Thus John and Mary know the same disjunctive truth, but Mary knows it a priori while John knows it a posteriori. Since the disjunction that either $289 + 365 = 654$ or there are cable cars in Switzerland inherits necessity from its first disjunct, John knows a necessary truth a posteriori.

The primary distinction between ways of knowing can be used to effect a secondary classification of things known. A truth p is a priori if and only if p can be known a priori. If we stipulate analogously that p is a posteriori if and only if p can be known a posteriori, then a truth may be both a priori and a posteriori, as with the disjunction that either $289 + 365 = 654$ or there are cable cars in Switzerland.

Alternatively, if we stipulate that a truth is a posteriori if and only if it is not a priori, then a truth that cannot be known a posteriori counts as a posteriori if it also cannot be known

a priori. There is no presumption that a truth can be known at all. Perhaps the best fit to current practice with the term is to stipulate that a truth is a posteriori if and only if it can be known a posteriori but cannot be known a priori. Even this way of drawing the distinction is subject to Kripke's forceful case (1980) that there are both contingent a priori truths and necessary a posteriori ones. Although some philosophers reject Kripke's arguments, even they now usually accept the burden of proof to show why the epistemological distinction between the a priori and the a posteriori should coincide with the metaphysical distinction between the necessary and the contingent.

To go further, we must clarify the terms 'independent' and 'experience'. One issue is that even paradigms of a priori knowledge depend in a sense on experience. For example, we supposedly know a priori that if it is sunny then it is sunny. But if our community had no direct or indirect experience of the sun or sunny weather, how could we understand what it is to be sunny, as we must if we are so much as to entertain the thought that it is sunny, let alone know that it is so?

The standard response is to distinguish between two roles that experience plays in cognition, one *evidential*, the other *enabling*. Experience is held to play an evidential role in our perceptual knowledge that it is sunny, but a merely enabling role in our knowledge that if it is sunny then it is sunny: we needed it only to acquire the concept *sunny* in the first place, not once we had done so to determine whether it currently applies. Experience provides our evidence that it is sunny, but not our evidence that if it is sunny then it is sunny; it merely enables us to raise the question. The idea is that an a priori way of knowing may depend on experience in its enabling role but must not depend on experience in its evidential role.

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

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

On further thought, however, that response causes more problems than it solves. For what prevents it from generalizing to outer experience? For example, part of the evidence that a massive comet or asteroid collided with the Earth about 250 million years ago is said to be that certain sediment samples from China and Japan contain certain clusters of carbon atoms. That those samples contained those clusters of atoms is a non-psychological fact. Of course, in some sense scientists' outer experience played a role in their access to the fact. But, by analogy with the logical and mathematical cases, the

relevant evidence is not the psychological process of undergoing those outer experiences, but rather the non-psychological physical facts to which that process enables us to have access. The role of the outer experience is purely enabling, not evidential. If so, what would usually be regarded as paradigm cases of a posteriori knowledge risk reclassification as a priori.

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

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

Of course, many attempts have been made to explicate the a priori – a posteriori distinction by introducing new theoretical apparatus. My aim here is not to discuss those attempts separately. Instead, I will address the distinction more directly, by comparing what would usually be regarded as a clear case of a priori knowledge with what would usually be regarded as a clear case of a posteriori knowledge. I will argue that the epistemological differences between the two cases are more superficial than they first

appear. The conclusion is not that the cases are borderline. I do not deny that they really are clear cases of a priori and a posteriori knowledge respectively, at least by bottom-up standards. In any case, showing that a distinction has borderline cases does not show that it is unhelpful for theoretical purposes. Rather, the appropriate conclusion is that the a priori – a posteriori distinction does not cut at the epistemological joints.

An analogy may be helpful with an argument of the same kind about a political distinction. If one aims to criticize the distinction between liberal and non-liberal policies, one achieves little by producing examples of policies that are neither clearly liberal nor clearly non-liberal. Every useful political distinction has borderline cases. But if one can produce an example of a clearly liberal policy that is politically only superficially different from a clearly non-liberal policy, then one has gone some way towards showing that the liberal – non-liberal distinction does not cut at the political joints.³

3. Here are two truths:

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

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

the predominant colour of their official cover). But let us describe two cases in more 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. He is not taught any rule like (1), connecting 'crimson' and 'red'. Through practice and feedback, he becomes very skilful in judging by eye whether something is crimson, and whether something is red. Now Norman is asked whether (1) holds. He has not previously considered any such question. Nevertheless, he can quite easily come to know (1), without looking at any crimson things to check whether they are red, or even remembering any crimson things to check whether they were red, or making any other new exercise of perception or memory of particular coloured things. Rather, he 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. That capacity is essential to much of our thinking, for instance when we reflectively assess conditionals in making contingency plans.⁴ No episodic memories of prior experiences, for example of crimson things, play any role. 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).

Naturally, that broad-brush description neglects many issues. For instance, what prevents Norman from imagining a peripheral shade of crimson? If one shade of crimson is red, it does not follow that all are. The relevant cognitive skills must be taken to include sensitivity to such matters. If normal speakers associate colour terms with central prototypes, as many psychologists believe, their use in the imaginative exercise may enhance its reliability. The proximity in colour space of prototypical crimson to prototypical red is one indicator, but does not suffice by itself, since it does not discriminate between ‘All crimson things are red’ (true) and ‘All red things are crimson’ (false). Various cognitive mechanisms can be postulated to do the job. We need not fill in the details, since for present purposes what matters is the overall picture. So far, we may accept it as a sketch of the cognitive processes underlying Norman’s a priori knowledge of (1).

Now compare the case of (2). Norman is as already described. He learns the complex phrase ‘recent volumes of *Who’s Who*’ by learning ‘recent’, ‘volume’, ‘*Who’s Who*’ and so on. He is not taught any rule like (2), connecting ‘recent volume of *Who’s Who*’ and ‘red’. Through practice and feedback, he becomes very skilful in judging by eye whether something is a recent volume of *Who’s Who* (by reading the title), and whether something is red. Now Norman is asked whether (2) holds. He has not previously considered any such question. Nevertheless, he can quite easily come to know (2), without looking at any recent volumes of *Who’s Who* to check whether they are red, or even remembering any recent volumes of *Who’s Who* to check whether they were red,

or any other new exercise of perception or memory. Rather, he assents to (2) after brief reflection along something like the following lines. First, Norman uses his skill in making visual judgments with ‘recent volume of *Who’s Who*’ to visually imagine a recent volume of *Who’s Who*. Then he uses his skill in making visual judgments with ‘red’ to judge, within the imaginative supposition, ‘It is red’. This involves the same general human capacity as before to transpose ‘online’ cognitive skills originally developed in perception into corresponding ‘offline’ cognitive skills subsequently applied in imagination. No episodic memories of prior experiences, for example of recent volumes of *Who’s Who*, play any role. As a result of the process, Norman accepts (2). Since his performance was sufficiently skilful, background conditions were normal, and so on, he thereby comes to know (2).

As before, the broad-brush description neglects many issues. For instance, what prevents Norman from imagining an untypical recent volume of *Who’s Who*? If one recent volume of *Who’s Who* is red, it does not follow that all are. The relevant cognitive skills must be taken to include sensitivity to such matters. As before, Norman must use his visual recognitional capacities offline in ways that respect untypical as well as typical cases. We may accept that as a sketch of the cognitive processes underlying Norman’s a posteriori knowledge of (2).

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

One response is to argue that at least one of the cases has been mislocated in relation to the a priori – a posteriori boundary. Perhaps Norman’s knowledge of (1) is really a posteriori, or his knowledge of (2) is really a priori (although presumably we did not make both mistakes). The risks of such a strategy are also obvious. If we reclassify Norman as knowing (1) a posteriori, we may have to do the same for all or most supposed cases of a priori knowledge, perhaps even of basic principles in logic and mathematics (such as standard axioms of set theory). For Norman’s knowledge of (1) did not initially seem atypical as a supposed case of a priori knowledge. On the other hand, if we reclassify Norman as knowing (2) a priori, we may still lose the distinction between a priori disciplines such as logic and mathematics and a posteriori disciplines such as physics and geography. Either way, we end up with an a priori – a posteriori distinction that cannot do much theoretical work.

Another response to the descriptions of how Norman knows (1) and (2) is more sceptical: it may be suggested that if the cognitive processes are really as described, then they are too unreliable to constitute genuine knowledge at all. However, this option is also unpromising for friends of the a priori – a posteriori distinction, for at least two reasons. First, it imposes an idealized epistemological standard for knowledge that human cognition cannot be expected to meet. None of our cognitive faculties is even close to being globally infallible. More local forms of reliability may suffice for knowledge; sceptics have not shown otherwise. Second, even if we are sceptical about knowledge in such cases, we should still assign belief in (1) and (2) some other sort of positive epistemic status, such as reasonableness, to which the a priori – a posteriori distinction should still apply in some form. Norman’s a priori reasonable belief in (1) and his a

posteriori reasonable belief in (2) could still be used in a similar way to argue against the depth of the new distinction. For purposes of argument, we may as well accept that Norman knows (1) and (2).

In the terms used in section 2, the question is whether Norman's experience plays an evidential or a merely enabling role in his knowledge of (1) and (2). Even in the case of (1), the role seems more than purely enabling. Consider Norbert, an otherwise competent native speaker of English who acquired the words 'crimson' and 'red' as colour terms in a fairly ordinary way, but has not had very much practice with feedback at classifying visually presented samples as 'crimson' or 'not crimson'. He usually makes the right calls when applying 'crimson' as well as 'red' online. By normal standards he is linguistically competent with both words. He grasps proposition (1). However, his inexperience with 'crimson' makes him less skilful than Norman in imagining a crimson sample. As a result, Norbert's reflection on whether crimson things are red comes to no definite conclusion, and he fails to know (1). Thus Norman's past experience did more than enable him to grasp proposition (1). It honed and calibrated his skills in applying the terms 'crimson' and 'red' to the point where he could carry out the imaginative exercise successfully. If Norman's experience plays a more than purely enabling role in his knowledge of (1), a fortiori it also plays a more than purely enabling role in his knowledge of (2).⁵

If the role of Norman's experience in his knowledge of (1) is more than purely enabling, is it strictly evidential? One interpretation of the example is that, although Norman's knowledge of (1) does not depend on episodic memory, and he may even lack all episodic memory of any relevant particular colour experiences, he nevertheless retains

from such experiences generic factual memories of what crimson things look like and of what red things look like, on which his knowledge of (1) depends. By contrast, Norbert fails to know (1) because his generic memories of what crimson things look like and of what red things look like are insufficiently clear. On this interpretation, Norman's colour experience plays an evidential role in his knowledge of (1), thereby making that knowledge a posteriori. But we have already seen that such reclassification is a risky strategy for defenders of the a priori – a posteriori distinction. Instead, it may be proposed, although colour experience can play an evidential role in a posteriori knowledge of what crimson things look like, and so indirectly in a posteriori knowledge of (1), we need not develop the example that way. The only residue of Norman's colour experience active in his knowledge of (1) may be his skill in recognizing and imagining colours.⁶ Such a role for experience, it may be held, is less than strictly evidential. Let us provisionally interpret the example the latter way. In section 5 we will reconsider, but reject, the idea that even supposed paradigms of a priori knowledge are really a posteriori.

Norman's knowledge of (2) can be envisaged in parallel to his knowledge of (1) as just envisaged. Although experience of recent volumes of *Who's Who* can play an evidential role in a posteriori knowledge of what such volumes look like, and so indirectly in a posteriori knowledge of (2), that is not what goes on with Norman. The only residue of his experience of recent volumes of *Who's Who* active in his knowledge of (2) is his skill in recognizing and imagining such volumes. That role for experience is less than strictly evidential. Nor does Norman's present experience play any more of an evidential role in his knowledge of (2) than it does in his knowledge of (1).

On this showing, the role of experience in both cases is more than purely enabling but less than strictly evidential. This reinforces a suspicion raised in section 2, that talk of ‘experience’ and ‘evidence’ does little to help us apply the a priori – a posteriori distinction top-down. Appeals to ‘observation’ as the hallmark of a posteriori knowledge hardly do better, for they leave us with the question: in what way does Norman’s knowledge of (2) involve observation while his knowledge of (1) does not?

The most salient difference between Norman’s knowledge of (2) and his knowledge of (1) is that (2) is contingent while (1) is necessary. That difference may be what inspires the idea that there must be a deep difference in his knowledge of them. But Kripke taught us not to read the epistemology of a truth off its metaphysical status. Of course, these two propositions do differ epistemologically:

(N1) It is necessary that crimson things are red.

(N2) It is necessary that recent volumes of *Who’s Who* are red.

For (N1) is a known truth, while (N2) is false, so not known, indeed presumably impossible, so unknowable. But an epistemological difference between (N1) and (N2) does not imply any epistemological difference between (1) and (2). In general, knowing a necessary truth does not imply knowing that it is necessary. For example, John in section 2 knew that either $289 + 365 = 654$ or there are cable cars in Switzerland by knowing that there are cable cars in Switzerland, without knowing that $289 + 365 = 654$; he did not know that it is necessary that either $289 + 365 = 654$ or there are cable cars in

Switzerland, even though it is indeed necessary. Likewise, Norman may know (1) without knowing (N1). He may be a sceptic about necessity, or never even entertain modal questions. In particular, Norman can know (1) without knowing (N1) and deducing (1) from it. Indeed, the idea that a precondition of knowing a necessary truth is knowing that it is necessary generates an infinite regress. For since (N1) is itself necessary, a precondition of knowing (N1) would be knowing (NN1), and so on ad infinitum:

(NN1) It is necessary that it is necessary that crimson things are red.

A subtler attempt to extract epistemological significance from the difference in modal status between (1) and (2) might exploit the modal nature of some proposed constraints on knowledge, such as various versions of reliability, sensitivity, and safety, many of which imply that one knows p only if falsely believing p is in some sense not too live a possibility.⁷ For since (1) is necessary, false belief in (1) is impossible, whereas false belief in (2) is possible, although not actual. Of course, since Norman is granted to know both (1) and (2), he satisfies any modal necessary condition for knowledge with respect to both truths. However, he might still be *more* reliable, or sensitive, or safe, with respect to (1) than to (2). But how could any such contrast make the difference between a priori and a posteriori knowledge? No constraint that all necessary truths trivially satisfy explains why some of them are known a priori, others only a posteriori.

To consider the point more fully, imagine Gull, who believes whatever his guru tells him. The guru tosses a coin to decide whether to assert to Gull Fermat's Last

Theorem (FLT), if it comes up heads, or its negation (\neg FLT), if it comes up tails. The coin comes up heads, and on the guru's testimony Gull obediently believes FLT, a necessary truth. It would generally be agreed that Gull does not know FLT. Indeed, in the non-technical senses of the terms, Gull's belief in FLT does not look particularly reliable, or safe, or sensitive to the facts. If the coin had come up tails, Gull would have wound up believing the necessary falsehood \neg FLT instead, in a parallel way. Any version of a reliability, sensitivity, or safety condition on knowledge non-trivially applicable to knowledge of a necessary truth p will concern possibilities of false belief in other propositions suitably related to p , not just of false belief in p itself. On such a dimension, we have no reason to expect Norman's knowledge of (1) to do better than his knowledge of (2). He may be just as prone to error in his judgments of colour inclusion as in his judgments of the colours of types of book; the distribution of errors in modal space may be much the same in the two cases. Once again, the difference in modal status between (1) and (2) is not what matters epistemologically.

The main effect of the modal difference between (1) and (2) may be to distract us from the epistemological similarity. The necessity of (1) prompts us to assimilate Norman's knowledge of (1) to stereotypes of a priori knowledge, which we can vaguely do because the role of experience is not strictly evidential. The contingency of (2) prompts us to assimilate his knowledge of (2) to stereotypes of a posteriori knowledge, which we can vaguely do because the role of experience is not purely enabling. Even having accepted Kripke's examples of the contingent a priori and the necessary a posteriori, we may still operate on the default assumptions that knowledge of necessary truths is a priori and that knowledge of contingent ones is a posteriori. Unlike Kripke's,

cases such as Norman's may trigger nothing to overturn either default, especially when they are considered separately, so by default we confidently classify knowledge of the necessary truth as a priori and knowledge of the contingent one as a posteriori, without noticing that there is no significant *epistemological* difference between them. We can use the qualifiers 'a priori' and 'a posteriori' that way if we like, but then we should not expect them to do much work in epistemology.

4. How widespread is the problem? It might be argued that although the a priori – a posteriori distinction does not mark any deep difference between Norman's knowledge of (1) and his knowledge of (2), the example is a special case, and that the distinction marks a deep difference in a wide range of other cases. A priori knowledge of logic and mathematics may be contrasted with Norman's knowledge of (1), and a posteriori knowledge by direct observation, preserved by memory, transmitted by testimony and extended by deductive, inductive, and abductive reasoning may be contrasted with his knowledge of (2). Thus, it might be claimed, the a priori – a posteriori distinction can still do plenty of useful work in epistemology after all.

I will argue that such an attitude is much too complacent. Many cases of a priori knowledge are relevantly similar to Norman's knowledge of (1), and many cases of a posteriori knowledge are relevantly similar to his knowledge of (2). Moreover, although epistemologists have become accustomed to treating the category of a priori knowledge as problematic, we still tend to treat the category of a posteriori knowledge as epistemologically explanatory. This attitude is particularly prevalent amongst those who deny that there is a priori knowledge. They think that it is clear enough how a posteriori

knowledge works, but hopelessly obscure how a priori knowledge could work.⁸ Once we appreciate how problematic the distinction itself is, we may rid ourselves of the illusion that we can understand what is going on in a case of knowledge by classifying it as a posteriori. The usual stereotype of a posteriori knowledge is just as epistemologically useless as the usual stereotype of a priori knowledge.

I will not discuss in detail how wide a range of other a posteriori knowledge resembles Norman's knowledge of (2). Since there is nothing very special about (2), it is fairly clear that if cognitive skills learnt online but applied offline can generate a posteriori knowledge of (2), without experience playing a strictly evidential role, then they can do likewise for many other truths. Examples include some knowledge of physical and practical possibility and of counterfactual conditionals.

In the case of (2), colour inclusions may look special. How much other putatively a priori knowledge resembles Norman's knowledge of (1)? He uses nothing like the formal proofs we associate with mathematical knowledge. A closer comparison is with knowledge of mathematical axioms, in particular with standard axioms of set theory (such as those of Zermelo-Fraenkel set theory).⁹

We may take as a typical example the Power Set Axiom, which says that every set has a power set, the set of all its subsets:

$$\text{PSA} \quad \forall x \exists y \forall z (z \in y \leftrightarrow z \subseteq x))$$

Here $z \subseteq x$ abbreviates the formula $\text{Set}(z) \ \& \ \text{Set}(x) \ \& \ \forall u (u \in z \rightarrow u \in x)$ ($\text{Set}(z)$ means: z is a set).¹⁰ Proofs of theorems throughout mathematics routinely and tacitly rely on PSA.

But how do mathematicians know that PSA is true? The problem here is not how mathematicians know that there are any sets at all, for if there are no sets then PSA is vacuously true (since both $z \in y$ and $z \subseteq x$ are always false). Rather, the problem is how mathematicians know that if there is a set, it has a power set.

Some textbooks motivate PSA in effect by telling readers that, unless they accept it, they will be unable to do set theory. They do not claim that accepting the axiom is necessary for understanding the language of set theory, in particular ‘Set’ and ‘ \in ’. They introduced those symbols at an earlier stage of the exposition. Once the axiom has been stated, readers are treated as grasping its content but potentially still wondering why to accept it; that is the point of the pragmatic motivation. It might be interpreted as an appeal to authority: take the author’s word for it, once you start working with this axiom you will see why it is needed in mathematics.

Other expositions of set theory attempt more intrinsic justifications of PSA. For instance:

If I have a set, then I can think of all possible subsets of this set. It is probably going to be a larger collection, but not so terribly much larger. It is reasonable to think of this as giving us back a set.¹¹

This is an implicit appeal to the principle of limitation of size, that things form a set if and only if there are not too many of them (fewer than absolutely all the things there are). Sometimes the appeal is backed up by the calculation that a finite set with just n members has just 2^n subsets. But of course PSA is intended to apply to infinite sets too.¹²

An alternative to limitation of size is the picture of sets as built up by an iterative process; at each stage one forms all possible sets of things already built up or given. This picture too is sometimes used to justify PSA:

[S]uppose x is formed at [stage] S . Since every member of x is formed before S , every subset of x is formed at S . Thus the set of all subsets of x can be formed at any stage after S .¹³

Of course, the apparently causal and temporal talk of forming sets at earlier or later stages is intended metaphorically, without commitment to any genuinely constructivist conception of sets. Nevertheless, the point of the metaphor is to appeal to the imagination, enabling us to think about the question in a more vivid, concrete, and perspicuous way, and in particular to convince us that there will be a stage after S , without which the power set never gets formed.¹⁴ The metaphor prompts us to undertake an imaginative exercise that makes offline use of our online skill in observing and engaging in processes of physical creation, a skill honed by past experience. This is not so distant from the imaginative exercise through which Norman came to know (1).

Something similar goes on in the justification from limitation of size. It starts with the supposition that 'I have a set', which already suggests a picture of the set as available to hand. On that supposition, 'I can think of all possible subsets of this set'. Of course, none of that is intended to suggest any idealist metaphysics of sets, on which it is essential to them to be thought by a subject. Rather, the aim is again to make us engage imaginatively with the question. The point of calling the subsets 'possible' is not to

emphasize that they could exist, for it is not in question that they actually do exist; it is to suggest that I could select them. Imagine that I have to hand three objects a , b , and c . They form a set from which I can make eight selections: the sets $\{a, b, c\}$; $\{a, b\}$; $\{a, c\}$; $\{b, c\}$; $\{a\}$; $\{b\}$; $\{c\}$; $\{\}$. They are the members of the power set of the set of the original three objects. My online experience of making different selections from amongst perceptually presented objects facilitates my offline imagined survey of all possible selections, and enables me to make the judgment in the quotation, ‘It is probably going to be a larger collection, but not so terribly much larger’. The cognitive tractability of the power set in such simple cases helps us accept PSA. Again, Norman’s knowledge of (1) is not so far away.

If standard axioms of set theory are justified by general conceptions of the sets, such as limitation of size or iterativeness, we may wonder how those general conceptions are in turn to be justified. Although the answer is hardly clear, all experience in the philosophy of set theory suggests that the attempt to make such a general conception of sets intuitively compelling must rest at least as heavily on appeals to the imagination with metaphors and pictures as do attempts to make intuitively compelling one of the standard set-theoretic axioms.

An alternative view is that such intrinsic justifications of set-theoretic axioms are secondary to extrinsic ones from their fruitfulness, their explanatory and unifying power. This need not involve Quine’s idea that mathematics is justified by its applications in natural science. Corresponding to the textbook’s implied injunction above to the mathematical novice, ‘wait and see’, applications in mathematics itself may be more relevant.¹⁵ Thus the strategy does not immediately commit one to an account of

mathematical knowledge as a posteriori, even though the envisaged abductive methodology is strongly reminiscent of the natural sciences.

Bertrand Russell describes a similar order of proceeding:

[I]nstead of asking what can be defined and deduced from what is assumed to begin with, we ask instead what more general ideas and principles can be found, in terms of which what was our starting-point can be defined or deduced.¹⁶

He observes:

The most obvious and easy things in mathematics are not those that come logically at the beginning; they are things that, from the point of view of logical deduction, come somewhere in the middle.¹⁷

This suggests knowledge of the ‘most obvious and easy things in mathematics’ as a better candidate than knowledge of the axioms of set theory to fit the stereotype of a priori knowledge. An example is knowledge that $2 + 2 = 4$, an arduously derived theorem of Russell and Whitehead’s system in *Principia Mathematica*. But if ordinary knowledge of elementary arithmetic is not by derivation from logically more basic principles, then presumably it is by something more like offline pattern recognition, and we still have not moved far from Norman’s knowledge of (1). Only the very lazy-minded could be content with the explanation that we know that $2 + 2 = 4$ ‘by intuition’. Even if it is true that we

do so in some sense of ‘intuition’, how does saying that constitute a genuine alternative to a view that assimilates our knowledge to Norman’s?

Even if experience plays no strictly evidential role in core mathematical practice, the suspicion remains that its role is more than purely enabling. Although we can insist that mathematical knowledge is a priori, it is unclear how it differs epistemologically from some examples of the a posteriori, such as Norman’s knowledge of (2).

Rather than pursuing the epistemology of mathematics further, let us see whether the stereotype of a priori fares any better in the epistemology of logic. For a simple example, consider the reflexivity of identity, the principle that everything is self-identical:

RI $\forall x x=x$

We are not asking about a priori knowledge that RI is a logical truth. We are just asking about a priori knowledge of RI itself, knowledge that everything is self-identical.

A tempting reaction is that anyone who doubts RI thereby just shows that they do not understand it. In the jargon, RI may be claimed to be *epistemologically analytic*. We need not discuss whether epistemological analyticity entails a priority, for it is false that basic logical truths are epistemologically analytic in the relevant sense. I know educated native speakers of English who deny that everything is self-identical, on the grounds that material substances that change their properties over time are not self-identical. I regard those speakers as confused, but the understanding they lack is primarily logical rather than semantic. Although they are mistaken about the logical consequences of identity, by

normal standards they are not linguistically incompetent with the English expressions ‘everything’, ‘is’, ‘self-’, and ‘identical’, and the way they are put together, nor with their counterparts in a formal first-order language with ‘=’. A language school is not the place for them to learn better. We might stipulate a sense of the tricky word ‘concept’ in which anyone who doubts RI counts as associating a different ‘concept’ with ‘=’ from any logically standard one, but recycling a theoretical disagreement as a difference in ‘concepts’ hardly clarifies the position.¹⁸

Although RI is not epistemologically analytic, it of course does not follow that it is not known a priori. One competent speaker may know a priori what another denies. But how is RI known? The universal generalization is unlikely to be an axiom of a formal system that develops innately in the human head. In a standard natural deduction system, RI is derived by the introduction rule $\forall I$ for the universal quantifier from a formula of the form $a=a$, which is itself a theorem (indeed, an axiom) by the introduction rule $=I$ for the identity sign.¹⁹ That is the formal analogue of imagining an object, within the scope of the imaginative supposition judging ‘it’ to be self-identical, and concluding that everything is self-identical. The $\forall I$ rule is subject to the restriction that the term a on which one universally generalizes must not occur (free) in any assumption from which the premise of the application of $\forall I$ was derived, otherwise one could derive $\forall x Fx$ from Fa . In this case the restriction is vacuously met; $a=a$ is a theorem and has no assumptions. Our informal thinking lacks a comparably clear way of keeping track of its assumptions. We make a judgment, perhaps within the scope of an imaginative supposition, but we may be unaware of its assumptions or sources. Thus it is often not transparent to us how far we can generalize. We may be imagining the case in a way that

is less generic or typical than we think. For example, those who deny (however mistakenly) that a changing thing is self-identical may charge that if we imagine an unchanging thing in evaluating RI we thereby beg the question in its favour.

Such reflections should not drive us into a general scepticism about our putative knowledge of universal generalizations. That would be an over-generalization of just the kind against which the reflections warn. They should not even drive us into a particular scepticism about our putative knowledge of RI. Knowing p does not require separately assessing in advance all possible fallacious objections to p . To require us to check that the imagined instance is typical of all members of a domain D before we universally generalize over D is to impose an infinite regress, for 'The imagined instance is typical of all members of D ' is itself a universal generalization over D . What matters for knowledge may be that we do safely imagine the instance in a relevantly generic way, even though the process is opaque to us. Surely we can quite easily know RI. Whether or not something changes is not really relevant to whether it is self-identical, so it does not matter whether we imagine a changing object or an unchanging one.

A resemblance between our knowledge of RI and Norman's knowledge of (1) is starting to emerge. Of course, an important aspect of Norman's knowledge of (1) is his offline imaginative use of capacities to apply colour terms calibrated perceptually online. Is there anything similar in our knowledge of RI? Experience involves a process of continually judging numerical identity or distinctness among objects perceived or remembered in a wide variety of guises. This cognitive capacity for judging identity and distinctness in experience is non-logical, for pure logic gives us only the barest formal constraints. If we have a non-logical capacity to make such identity judgments, we need

no additional logical capacity corresponding to the rule =I to make identity judgments of the special syntactic form $a=a$. After all, we could use the non-logical capacity to judge $a=b$ and $b=a$ (for some suitable term b syntactically distinct from a) and then apply the transitivity of identity to deduce $a=a$. The transitivity of identity does not depend on its reflexivity.²⁰ A simpler and more plausible way of using a non-logical capacity to make judgments of identity and distinctness to judge $a=a$ would be directly to feed in the term a twice over as both inputs to some device for comparison, which would trivially return a positive result. That can be done online or offline.

Even the trivial comparison in $a=a$ can be mishandled. If the name a denotes an enduring, changing substance, but one associates the first token of a with the properties the object had at a time t (attributed in the present tense, not relativized to t) and the second token of a with the properties the object had at another time t^* (also attributed in the present tense, not relativized to t^* , and incompatible with the former properties), then the output from the identity test is a false negative. Through experience of material things undergoing slow large changes, one becomes less prone to such mistakes, although some adult metaphysicians still manage to make them, and so deny RI.

The foregoing remarks are not intended to suggest that knowledge of RI is a posteriori. Classify it as a priori by all means, but do not let that blind you to how much it has in common with a posteriori knowledge of identity and distinctness, just as Norman's a priori knowledge of (1) has so much in common with his a posteriori knowledge of (2).

The salient difference between (1) and (2) is modal rather than epistemological: (1) is metaphysically necessary, (2) metaphysically contingent. By contrast, $a=a$ and $a=b$ do not differ in modal status if both are true and the terms a and b are proper names or

other rigid designators, for an identity claim with such terms is metaphysically necessary if true at all.²¹ But in that case the salient difference between the formulas $a=a$ and $a=b$ is logical rather than epistemological. For $a=a$ but not $a=b$ is a logical truth. The property of logical truth is not demarcated epistemologically but by more formal criteria.²² Just as the modal difference between (1) and (2) makes us overestimate the strictly epistemological difference between Norman's knowledge of (1) and his knowledge of (2), so the logical difference between $a=a$ and $a=b$ makes us overestimate the strictly epistemological difference between our knowledge of $a=a$ and our knowledge of $a=b$.

Naturally, far more work would have to be done to confirm the foregoing hints about logical and mathematical knowledge. Nevertheless, the indications so far suggest that what are often counted as the clearest cases of a priori knowledge are much less different epistemologically than they are usually depicted as being from cases of a posteriori knowledge. The epistemological similarity of Norman's a priori knowledge of (1) to his a posteriori knowledge of (2) is no isolated case. The usual stereotype of a priori knowledge is seriously misleading, because it omits a pervasive role of experience that is more than purely enabling, although less than strictly evidential.

5. The inadequacy of the usual stereotype of a priori knowledge may seem to support the idea that we can make progress by following some followers of Quine in classifying all knowledge as a posteriori. Doing so would at least have the negative advantage of not putting a distinction where there is no deep difference. But does it also yield some positive understanding of the general nature of knowledge?

On Quine's picture, a theory faces the tribunal of experience collectively, not sentence by sentence. Taken at face value, the image implies that two consequences of a theory cannot differ in epistemic status. But that is absurd. For Quine, the totality of a person's beliefs constitute a theory (perhaps an inconsistent one), their total theory of the world, but who thinks that two of a person's beliefs cannot differ in epistemic status? If some of my beliefs constitute knowledge, it does not follow that all of them do; it does not even follow that all of them are true. One's beliefs about science and mathematics may be on average epistemically better off than one's beliefs about religion and politics, or vice versa. Indeed, experience favours the belief that experience favours some beliefs more than others. But Quine's holism does not justify restricting his maxim to local theories rather than global ones. He is surely right to at least this extent: no two of our beliefs are in principle epistemically insulated from each other.

To make progress, we need a more developed model, on which an individual belief has its own epistemic status, but that status depends in principle on the epistemic status of each other belief. Holism is far more plausible as a claim about the pervasive interdependence of epistemic status than as the claim that only whole theories have epistemic status. The obvious and best-developed candidate for such a model is some form of Bayesian epistemology. It assigns evidential probabilities to individual propositions, subject to standard axioms of probability theory, which constrain the overall distribution of probabilities to all propositions.²⁴ The paradigmatic way of updating evidential probabilities is by conditionalization on new evidence, encapsulated in a proposition e . The new evidential probability $\text{Prob}_{\text{new}}(p)$ of any proposition p is the old conditional evidential probability $\text{Prob}_{\text{old}}(p | e)$ of p on e , which is equal to the ratio

$\text{Prob}_{\text{old}}(p \ \& \ e) / \text{Prob}_{\text{old}}(e)$ whenever $\text{Prob}_{\text{old}}(e) > 0$. Conditionalization is a global process; one overall probability distribution, Prob_{new} , replaces another, Prob_{old} .

However, Bayesian epistemology does not vindicate Quine's rejection of the a priori. For standard axioms of probability theory constrain every probability distribution to assign probability 1 to any theorem of classical propositional logic, and probability 0 to its negation. Probabilistic updating on new evidence cannot raise or lower the probability of theorems or anti-theorems. That is not just an optional convention. Loosening it deprives probability theory of the mathematical structure on which its utility depends. Although minor concessions to specific non-classical logics may not destroy that utility entirely, any version of probability theory worth having will give such a privileged status to some core of logic.

In Bayesian epistemology, logical truths are not the only propositions to enjoy a good epistemic status that they do not owe to the evidence. Let e conjoin all the relevant evidence, and Prob_{new} be the result of conditionalizing Prob_{old} on e as above. We may assume that the evidence was not certain in advance; $\text{Prob}_{\text{old}}(e) < 1$. Suppose that a hypothesis h is well supported by e , so $\text{Prob}_{\text{new}}(h)$ is high. Consider the material conditional $e \rightarrow h$. Since it is a logical consequence of h , $\text{Prob}_{\text{new}}(e \rightarrow h)$ is at least as high as $\text{Prob}_{\text{new}}(h)$. But we can prove that either $\text{Prob}_{\text{new}}(e \rightarrow h) = \text{Prob}_{\text{old}}(e \rightarrow h) = 1$ or $\text{Prob}_{\text{new}}(e \rightarrow h) < \text{Prob}_{\text{old}}(e \rightarrow h)$.²⁵ In other words, either $e \rightarrow h$ was already certain prior to the evidence, which did not confirm $e \rightarrow h$, or the evidence disconfirmed $e \rightarrow h$. Thus $e \rightarrow h$ enjoys a good epistemic status, because $\text{Prob}_{\text{new}}(e \rightarrow h)$ is high, but it does so despite the evidence or independently of it.

When holistic epistemology is made rigorous, the results do not support the idea that the only way of enjoying high epistemic status is by confirmation through experience; they do the opposite. That is not to deny the strong similarities between the epistemology of logic and the epistemology of other sciences evident in debates over proposals to revise or extend classical logic. For we cannot assume that those similarities are properly articulated on the model of confirmation or disconfirmation through experience.

Rather than appealing to formal models, those who claim that all knowledge is empirical or a posteriori may suggest that we understand the paradigm of such knowledge, simple cases of observational knowledge, well enough for the informal proposal to assimilate all knowledge to the paradigm to be illuminating and non-trivial. Although there are manifest differences between the paradigm and cases of highly theoretical knowledge, the idea is that on sufficiently deep analysis they will turn out to be differences in complexity, not in fundamental nature.

How well do we understand the paradigm, simple cases of observational knowledge? Presumably, the picture is that in such cases sense perception is a channel for a causal connection between the truth of a proposition p and the agent a 's belief in p , creating a strongly positive local correlation between truth and belief. We may symbolize the correlation as $p \Leftrightarrow B_a p$. The proposition p should paradigmatically concern the state of the environment, not the state of the agent, for otherwise the case is too special to be a suitable model for knowledge in general.

A first step in making the model less simplistic is to note that the correlation depends on the receptivity of the agent. If a is too far from the relevant events or shuts

her eyes or has bad eyesight, the preconditions for the correlation may not be met. We can symbolize this one-way correlation between receptivity and the previous two-way correlation as $R_ap \Rightarrow (p \Leftrightarrow B_ap)$.²⁶ Even at this utterly elementary level, it is clear that a causal connection between truth and belief is not the only way to achieve such a set-up. For suppose that p is a necessary truth (as it were, $\Rightarrow p$), and that the receptivity of the agent by itself causes the belief ($R_ap \Rightarrow B_ap$). Then we have $R_ap \Rightarrow (p \Leftrightarrow B_ap)$, even though there is no causal connection between p and B_ap . The receptivity condition R_ap here should not be envisaged as some mystical state of opening one's soul to Platonic heaven; it may be a mundane psychological process, for example of calculation. Thus the $R_ap \Rightarrow (p \Leftrightarrow B_ap)$ model carries no commitment to conceiving the modelled epistemic states as all a posteriori. For all it implies, some of them are a priori. Perhaps surprisingly, treating simple observation as the paradigm to which all knowledge must be assimilated does not in principle commit one to a uniformly a posteriori epistemology.

As an alternative paradigm of knowledge, many self-described naturalists prefer the experimentally-based findings of the natural sciences. A conception of all knowledge as a posteriori or empirical may be an attempt to assimilate it all to natural science. Of course, one of the most salient obstacles to any such attempt is mathematics. Obviously, theorems of mathematics do not normally have direct experimental support. For Quine, they have indirect experimental support because mathematics is part of our total scientific theory, which is confirmed as a whole (if at all) by experimental tests. But we have already seen that theories are not the only bearers of epistemic status. The practice of the natural sciences themselves requires evaluating the epistemic status of much smaller units: for example, should we believe a report of a given astronomical observation or

experimental result? Once we ask more discriminating questions about the epistemic status of individual axioms and theorems of mathematics, it becomes much harder to tell a plausible story on which they owe that status primarily to experimental support.

Although some axioms and theorems are in a better epistemic position than others, that has far more to do with considerations internal to mathematics than with experimental support.²⁷ The same holds even more obviously for axioms and theorems of logic.

On present evidence, the slogan ‘All knowledge is a posteriori’ or ‘All knowledge is empirical’ is defensible only if the term ‘a posteriori’ or ‘empirical’ is emptied of all serious content. Unfortunately, that does not deprive the underlying prejudice of influence. Like other prejudices, it acts selectively, for instance by imposing more severe demands for external justification on armchair methods in philosophy than on other methods of inquiry.

6. We should not react to the inadequacy of the usual stereotype of a priori knowledge by declaring all knowledge a posteriori. Conversely, of course, we should not react to the inadequacy of the usual stereotype of a posteriori knowledge by declaring all knowledge a priori. Since the terms ‘a priori’ and ‘a posteriori’ are not meaningless by normal standards, some difference between a priori and a posteriori knowledge remains. But that does not rehabilitate the distinction as of great theoretical value for epistemology. After all, there is a difference between plants that are bushes and plants that are not bushes, but it does not follow that that distinction is of great theoretical value for botany.²⁸ When we start investigating some phenomena, we have little choice but to classify them according to manifest similarities and differences. As our understanding deepens, we may recognize

the need to reclassify the phenomena on less obvious dimensions of greater explanatory significance. Distinctions that aided progress in the early stages may hinder it later on. The a priori – a posteriori distinction is a case in point.

Cognitive psychology will have much to offer the epistemologist's attempt to overcome philosophical prejudices and classify according to deeper and less obvious similarities and differences. But that does not mean a reduction of epistemology to cognitive psychology. Epistemological questions are typically at a higher level of generality than those of cognitive psychology: for example, they may concern all knowledge, all epistemic probability, or all rationality. Epistemology also engages more fully with evaluative questions about how knowledge is better than mere true belief, or rationality than irrationality, or whether we ought to proportion our beliefs to the evidence. Of course, some 'naturalized' epistemologists abjure such evaluative questions as 'unscientific'. In doing so, they seem more under the influence of logical positivism, or the prejudice discussed in section 5, than of actual scientific practice. In particular, cognitive psychologists do not abjure all evaluative judgments about what is rational or irrational, in their experimental studies of irrationality. Moreover, pursuing any scientific inquiry involves making numerous judgments that are evaluative in the way characteristic of epistemology. Are these data trustworthy? Is that argument valid? Which of these theories is better supported by the evidence? Even if those questions play a merely instrumental role in the natural sciences, there is no warrant for the idea that scientific practice somehow discredits the attempt to inquire more generally into the nature of phenomena like trustworthiness, validity, and evidential support. Indeed, it seems

contrary to the scientific spirit to disapprove of systematic general inquiry into such matters while nevertheless continually relying on judgments about them.

Beyond the potential contribution of cognitive psychology, we need to develop a more detailed, precise and specifically epistemological vocabulary for describing the fine structure of examples such as those in sections 3 and 4, involving the offline application in the imagination of cognitive skills originally developed in perception, especially when they involve generic imaginary instances used to reach general non-imaginary conclusions. That is likely to prove a far more fruitful project for epistemology than yet another attempt to reconstruct the tired-out distinction between the a priori and the a posteriori or to stretch the latter thin enough to cover all knowledge.

Notes

1 This paper develops the brief critique of the distinction between a priori and a posteriori knowledge in Williamson 2007, pp. 165-9. Earlier versions of the material were presented at King's College London, the University of Santiago de Compostela, Moscow State University, St Petersburg State University, the University of Hertfordshire, and a course on Mind, World, and Action at the Inter-University Centre in Dubrovnik. I am grateful to audiences there and to Anna-Sara Malmgren for questions and discussion.

2 Many contemporary epistemologists, especially those of an internalist bent, treat the distinction as primarily a classification of forms of justification rather than of knowledge. For reasons explained in Williamson 2000, I treat the classification of forms of knowledge as primary. Friends of justification should not find much difficulty in reworking the arguments of this paper in their terms.

3 I am not actually endorsing the conclusion that the liberal – non-liberal distinction does not cut at the political joints, because I have not actually given an example of a clearly liberal policy that is politically only superficially different from a clearly non-liberal policy.

4 See Williamson 2007, pp. 137-78.

5 It is even more implausible that experience plays a purely enabling role in the more complex examples in Williamson 2007, pp. 166-7; (1) is used here for its simplicity.

6 Compare the huge debate on Mary's (a posteriori?) knowledge of what red things look like (Jackson 1982).

7 For some discussion see Williamson 2000, pp. 123-30 and 147-63.

8 For an example of this attitude see Devitt 2005.

9 The axioms of group theory are simply clauses in the mathematical definition of 'group' and so raise no distinctive epistemological problem; likewise for the axioms for other kinds of algebraic or geometrical structure. The primary role of the axioms of set theory in mathematics is quite different. Proofs in all branches of mathematics rely on their truth. They are not clauses in the mathematical definition of 'set' (there is no such definition analogous to the mathematical definition of 'group'). Although they can be adapted to serve as such clauses in a mathematical definition of 'cumulative hierarchy' or the like, that is a secondary use; even proofs about all cumulative hierarchies rely on the axioms of set theory in their primary role. Some globally structuralist philosophers of mathematics may urge a different attitude, but mathematicians have not yet seen fit to indulge them, nor is it clear how they could coherently adopt a globally structuralist attitude.

10 The quantifiers in PSA are not restricted to sets. In applying mathematics, we may need the power set of a set of concrete objects. If x is a non-set, it has no subsets, so y can be $\{\}$ (if we want a set) or x itself (if we want a non-set). If x is a set, then $\{\} \subseteq x$, so y must be a set since $\{\} \in y$ and only sets have members. If we did not impose $\text{Set}(z)$ as a condition on $z \subseteq x$, every non-set would vacuously count as a subset of any set x and so have to belong to the power set of x .

11 Crossley, Ash, Brickhill, Stillwell, and Williams 1972, p. 61.

12 An infinite set with κ members also has 2^κ subsets, but in the infinite case the definition of 2^κ depends on PSA rather than offering it independent support. For the history of the limitation of size principle and scepticism about its capacity to motivate PSA see Hallett 1984.

13 Shoenfield 1977, p. 326. For a more philosophical account of the iterative conception of set see Boolos 1971.

14 The appeal to the imagination is explicit at pp. 323-4 of Shoenfield 1977, in his general account of the process of set formation. Note that he is writing as a mathematician for mathematicians and others, not as a philosopher.

15 See Maddy 2011.

16 Russell 1919, p. 1. According to Russell, the latter order ‘characterises mathematical philosophy as opposed to ordinary mathematics’. If Maddy is right, it also characterises set theory as a branch of ordinary mathematics.

17 Russell 1919, p. 2.

18 See Williamson 2007, pp. 73-133, for a more detailed critique of epistemological analyticity.

19 Matters are more complicated in free logic, where the term a is not guaranteed to denote anything in the domain of quantification.

20 The standard proof of transitivity uses the elimination rule $=E$, the indiscernibility of identicals, but not the introduction rule $=I$.

21 The classic defence of the necessity of identity is, of course, Kripke 1980. On some metaphysical views, what is necessary in the case of both identities is only that if the objects exist then they are identical.

22 For present purposes logical truths are treated as formulas rather than propositions. The truth of $a=b$ does not make it the same logical truth as $a=a$.

23 ‘It is overwhelmingly plausible that *some* knowledge is empirical, “justified by experience.” The attractive thesis of naturalism is that *all* knowledge is; there is only one way of knowing.’ (Devitt 2005, p. 105).

24 More accurately, the constraints are on a distribution of probabilities to all propositions in the σ -field of propositions that receive probabilities at all.

25 $\text{Prob}_{\text{old}}(e \ \& \ \neg h) = \text{Prob}_{\text{old}}(e \ \& \ \neg h \ | \ e)\text{Prob}_{\text{old}}(e) + \text{Prob}_{\text{old}}(e \ \& \ \neg h \ | \ \neg e)\text{Prob}_{\text{old}}(\neg e)$.

But $\text{Prob}_{\text{old}}(e \ \& \ \neg h \ | \ \neg e) = 0$ and $\text{Prob}_{\text{old}}(e \ \& \ \neg h \ | \ e) = \text{Prob}_{\text{new}}(e \ \& \ \neg h)$, so

$\text{Prob}_{\text{old}}(e \ \& \ \neg h) = \text{Prob}_{\text{new}}(e \ \& \ \neg h)\text{Prob}_{\text{old}}(e)$. Since $\text{Prob}_{\text{old}}(e) < 1$,

$\text{Prob}_{\text{old}}(e \ \& \ \neg h) < \text{Prob}_{\text{new}}(e \ \& \ \neg h)$ unless $\text{Prob}_{\text{new}}(e \ \& \ \neg h) = 0$. In the former case,

$\text{Prob}_{\text{new}}(e \ \rightarrow \ h) = 1 - \text{Prob}_{\text{new}}(e \ \& \ \neg h) < 1 - \text{Prob}_{\text{old}}(e \ \& \ \neg h) = \text{Prob}_{\text{old}}(e \ \rightarrow \ h)$. In the

latter case, $\text{Prob}_{\text{old}}(e \ \& \ \neg h) = 0$ too, so $\text{Prob}_{\text{old}}(e \ \rightarrow \ h) = 1$.

26 The correlation between receptivity and the former two-way correlation may itself be two-way ($R_{ap} \Leftrightarrow (p \Leftrightarrow B_{ap})$), if receptivity is the only way of setting up the first correlation, but the other direction does not matter for present purposes.

27 For example, in the set theory ZFC, the Axiom of Replacement is usually considered to be better established than the Axiom of Choice, but not because it has more experimental support. For other objections to a purely holistic account of the confirmation of mathematics as part of total science see Maddy 2007, pp. 314-17.

28 We are not concerned with more radical redefinitions of the words in the sentence.

29 For an introduction to psychological work on the cognitive value of the imagination, the case specifically relevant to the examples in section 3, see Harris 2000.

References

- Barwise, J. (ed.) 1977. *Handbook of Mathematical Logic*. Amsterdam: North-Holland.
- Boolos, G. 1971. 'The iterative conception of set', *The Journal of Philosophy* 68, pp. 215-32.
- Crossley, J.N., Ash, C.J., Brickhill, C.J., Stillwell, J.C., and Williams, N.H. 1972. *What is Mathematical Logic?* Oxford: Oxford University Press.
- Devitt, M. 2005. 'There is no *a priori*', in Steup and Sosa 2005, pp. 105-15.
- Hallett, M. 1984. *Cantorian Set Theory and Limitation of Size*. Oxford: Clarendon Press.
- Harris, P.L. 2000. *The Work of the Imagination*. Oxford: Blackwell.
- Jackson, F. 1982. 'Epiphenomenal qualia', *The Philosophical Quarterly*, 32, pp. 127-36.
- Kripke, S. 1980. *Naming and Necessity*. Oxford: Blackwell.
- Maddy, P. 2007. *Second Philosophy: A Naturalistic Method*. Oxford: Oxford University Press.
- Maddy, P. 2011. *Defending the Axioms: On the Philosophical Foundations of Set Theory*. Oxford: Oxford University Press.
- Quine, W.V.O. 1951. 'Two dogmas of empiricism', *Philosophical Review*, 60, pp. 20-43.
- Russell, B.A.W. 1919. *Introduction to Mathematical Philosophy*. London: George Allen and Unwin.
- Shoenfield, J.R. 1977. 'Axioms of set theory', in Barwise 1977, pp. 321-44.
- Steup, M., and Sosa, E. (eds.) 2005. *Contemporary Debates in Epistemology*. Oxford: Blackwell.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, T. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.