

A Priority and Externalism¹

#0 Introduction

If an epistemological distinction fails to carve at the epistemological joints, then it is not worthy of serious and protracted discussion. The residual issue of whether the putative distinction is incoherent or merely gerrymandered ought not to strike anyone as especially important. My own externalist commitments – epistemological and semantic -- 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.² I shall be defending that perspective in what follows. Discussions of a priority typically assume that there is a distinguished subclass of beliefs whose epistemic status does not depend upon experiential encounter with the world. Yet the project of delimiting such a class in a non-gerrymandered fashion turns out to be surprisingly difficult. As an organizing theme, I shall examine the bearing of a safety based account of knowledge on the traditional conception of a priori knowledge.

#1 Environment Independence

Epistemologists divide according to whether they take knowledge or, instead, some kind of justification (or ‘warrant’ or ‘rationality’ or ‘entitlement’) as the starting point for foundational inquiry. In line with my preferred orientation, much of the discussion that follows will concern a priori knowledge. In a final section, I shall speak to the topic of a priori justification.

Let us begin with a simple externalist picture of knowledge, a version of the safety account:³

S knows p iff there is no close world where S makes a mistake that is relevantly similar to his actual belief that p.⁴

That there is a close world where S is mistaken about p does not show that he does not actually know p, since he may use a relevantly different method at that close world. (I could easily have asked an unreliable informant the score of the game rather than checked the scoreboard, and at a close world where the score is different the informant provides me with the score that obtains at the actual world.) And that there is no close world where

¹ I am grateful to Elizabeth Camp, Maya Eddon and Sanford Goldberg for comments on an earlier draft, and to Timothy Williamson and an audience at Oxford for helpful discussion. I am especially grateful to David Manley for discussions that helped to develop and sharpen the material in this paper. We hope to develop some of these themes further in *Thoughts and their Objects*, forthcoming.

² A similar conclusion is argued for in Sarah Sawyer, ‘The Epistemic Divide,’ *Southern Journal of Philosophy*, 2001, 385-401.

³ Versions of the safety idea can be found, inter alia, in Ernest Sosa, ‘How Is Knowledge Modally Related to What is Known?’ in *Philosophical Topics* 26, 1999, 373-384, and Timothy Williamson, *Knowledge and Its Limits*, Oxford University Press, 2000. The formulation here is different.

⁴ The discussion of empty thoughts below indicates that we need to be liberal about the notion of ‘mistake’.

S is mistaken about p does not show that he actually does know that p. For there may be a close world where S is not mistaken about P but does make a mistake that is sufficiently similar to the actual case in relevant respects. (Suppose I form a demonstrative belief of Jones that he is Jones but could easily have formed a demonstrative belief of Jones' Twin – who was also in the room -- that he is Jones. The possible mistake concerns a different proposition but still undermines actual knowledge.)⁵ This point is relevant to evaluating beliefs in mathematical necessities too. Suppose I use an unreliable method to form a belief in a mathematical necessity, P. Even though I could not easily have believed P falsely, there are relevantly similar cases in which I am in error.

My interest here is not in whether the picture allows us to make sense of mathematical knowledge: certainly it does. It is rather to ask what a priority comes to when knowledge is conceived along these lines. It is often that thought that in a case of a priori knowledge, the status of a belief as knowledge does not constitutively depend on the external environment (this being one natural take on the idea that a priori knowledge is independent of perceptual experience). Let us pursue this idea.⁶

Here is a crude first pass:

x knows p a priori after duration d of x's existence iff any possible intrinsic duplicate y of x knows p after duration d of y's existence.⁷

The truth of semantic externalism prohibits us from this way of articulating the idea. My semantic life does not supervene on my inner life. As a result, there may be cases where an intrinsic duplicate does not think p at all, let alone know p. Suppose for example, that the semantic value of 'thin' in my mouth constitutively depends on the dispositions of my linguistic community. Suppose at t I know that thin people are thin. Some possible duplicate of me does not think about that proposition at all, but instead, believes a proposition about thin+ things (where thin+ is the semantic value of 'thin' in his mouth). Even more obviously, while I may be reckoned to know a priori that if I am very happy

⁵ The picture is not intended as a way of adjudicating cases where intuitions vary about whether someone knows. At best it offers a kind of diagnosis of such disputes: intuitions will vary as to which kinds of possible mistakes are relevantly similar, and coordinately, which kind of individuation of methods is relevant to the case at hand. Insofar as knowledge is conceptually basic, no specification of relevant similarity conceptually prior to knowledge will be available.

⁶ One finds something like this idea in Philip Kitcher's well known attempt to capture the traditional conception of the a priori ('A Priori Knowledge,' *The Philosophical Review*, Vol. 89, No 1, Jan 1980, 3-23.) Here is his account: (i) Someone knows p a priori iff X knows p and X's belief was produced by a process which is an a priori warrant for it. (ii) α is an a priori warrant for X's belief that p if and only if α is a process such that, given any life e sufficient for X for p, then (a) some process of the same type could produce in X a belief that p (b) if a process of the same type were to produce in X a belief that p then it would warrant X in believing that p (c) if a process of the same type were to produce in X a belief that p then p. (pp. 9-10). Note that 'a life sufficient for X for p' means 'X could have had that life and gained sufficient understanding to believe that p.' (pp. 5-6) There is plenty to worry about here. It does not handle cases where experience provides knowledge destroying counterevidence (the experience includes experts who all say 'Not-P'). And it does not take stock of the fact that a life sufficient for the conclusion of a paradigmatically a priori proof might not be a life sufficient for one or more of the premises. Standing back from these matters of detail, though, it is quite clear that Kitcher's basic idea is that a process warrants a belief a priori iff, no matter how the environment is, that process is a warrant provider. The considerations that follow will make trouble for that vision, no matter what the details of its implementation.

⁷ We require here that the possible individual duplicate the entire intrinsic history of x.

then I am happy, a possible intrinsic duplicate who is numerically distinct from me will not be thinking *that* proposition, since it is about me, not the duplicate. Here is a more promising way of capturing the intuition of environment independence:

Independence₁: A case of x's believing p is a case of a priori knowledge iff for any possible intrinsic duplicate y⁸, the counterpart in y of x's belief that p is a case of knowledge.⁹

The notion of a counterpart is intuitive enough here. Crucially, a counterpart of my knowledge that thin people are thin might be knowledge of a different proposition. Further qualifications will be needed if one is the kind of semantic externalist who thinks that there are possible intrinsic duplicates of oneself who do not have any beliefs at all, or who thinks that there are at least some beliefs such that there is a possible intrinsic duplicate with no counterpart belief. Such a theorist might consider the following, alternative proposal:

Independence₂: A case of x's believing p is a case of a priori knowledge iff for any possible intrinsic duplicate y such that the counterpart in y of x's belief that p is a case of believing that p, the counterpart in y of x's belief is a case of knowledge.^{10 11}

(This is already a kind of retreat. The idea behind Independence₁ is that the environment makes no difference to *whether* an episode is a case of knowledge – it only makes a difference to *which* proposition one knows. Once we take seriously the possibility of a duplicate's having empty thoughts, the contribution of the environment has to be recognized as more substantial.)

Let us try to embed these ideas within a safety theoretic conception of knowledge. In a posteriori knowledge, the environment plays the role of providing a *safe haven*. Take that part of an a posteriori knowledge-delivering process that is intrinsic to the subject. That intrinsic profile can be embedded in environments where the subject's counterpart belief is not knowledge at all. But according to the gloss encoded by Independence₁, a priori knowledge is distinctive in that it stands in no special need of a safe haven from the environment, since the environment can't but provide a safe haven: the intrinsic character

⁸ Some philosophers will think that my intrinsic nature radically undetermines the laws governing me. On such a view, some possible intrinsic duplicates might have radically different dispositions, and enjoy radically different causal connections between even the components of its own inner life. Such a philosopher had better restrict the thesis to intrinsic and nomic twins.

⁹ Here I take intrinsic duplication to include both my physical and conscious life beneath the skin. Versions of the principle that merely include conscious life will be even more easily susceptible to troubling examples.

¹⁰ It will not do to say 'counterpart who believes that p'. Suppose x believes P on good grounds and P' on bad grounds. Owing to externalist considerations, there may be cases where a duplicate is such that the counterpart of x's belief that p is a belief that p' and the counterpart of x's belief that p' is a belief that p. In that case the duplicate would believe that p but not know p. But it doesn't seem that this should all by itself vitiate a claim of a priority for the actual belief that p.

¹¹ I don't mean to suggest that the two proposals are the only reasonable candidates in the vicinity. One gets a conception intermediate between the two in the text by substituting 'a case of believing some proposition' for 'a case of believing that p' in the second conception. The considerations that follow will apply to this conception as well.

of the process guarantees knowledge, whatever the vicissitudes of the environment. Meanwhile, Independence₂ is motivated by the following idea. In a case of a posteriori knowledge that P, the environment plays a dual role: it affords the connections to the world that allow to think to have the propositional object P as the object of belief in that case, and it provides a safe haven for that belief. In a case of a priori knowledge that P, by contrast, the first role is sufficient to ensure that certain inner lives are cases of knowledge.

How does the environment independence picture stand up? The trouble is that, with a bit of imagination, it is not hard to contrive ways that even paradigmatically a priori beliefs stand in need of a safe haven. Let me offer a few illustrative examples, the second and third of which turn on semantic externalism.

A The Danger of Bad Influence

We are all familiar with fake barn cases.¹² A person sees a barn in an environment where there are lots of non-barns that look just like barns. Even though he forms a true perceptual belief of the barn that it is a barn, many of us are inclined not to count the belief as knowledge. The safety based approach that I have sketched offers a natural diagnosis of that judgment. Supposing barns are essentially barns, the proposition believed is a necessary truth. Still the belief is not knowledge, since mistakes that are relevantly similar populate close worlds. The same judgment would no doubt be elicited if the real barn was in the vicinity of barn gas that induces the hallucination of barns. But the last case can easily be adapted to allegedly a priori beliefs. Suppose there exists a priori gas that induces the phenomenology of blatant obviousness for false propositions. Consider a person who believes a proposition not for any empirical reason but because the phenomenology of obviousness causes him to do so. Suppose the claim in question is that all bachelors are men. Consider a duplicate of that person who is embedded in an environment riddled with a priori gas. As a matter of luck he does not stumble into the gas. He in fact forms the belief that all bachelors are men. But he could very easily have stumbled into the gas and believed – due to felt obviousness --- that all bachelors are women.¹³ Insofar as one judges that the person does not know in fake barn cases, it is natural enough to judge that the person does not know in a priori gas cases. But this means that if cling to the environment dependence idea, very few of our beliefs will count as a priori. Even so-called a priori beliefs need a safe haven.

The example turns on close worlds where one comes under a bad influence. Simple examples abound. The bad influence could be a priori gas. It could be a neuron-disturbing magnetic field or some belief disturbing quantum entanglement. More mundanely, it could be a group of persuasive but sophisticated intellectuals that could very easily have induced different habits and tastes in one's abstract belief formation (suppose that whether one believed that arithmetic was true or a useful fiction depended on which

¹² See, notably, Alvin Goldman, 'Discrimination and Perceptual Knowledge,' *Journal of Philosophy*, November 1976, 771-791. For discussion of the delicate nature of our intuitions in these cases, see Tamar Gendler and John Hawthorne, 'The Real Guide to Fake Barns: A Catalogue of Gifts for your Epistemic Enemies,' *Philosophical Studies*, June 2005, 331-352.

¹³ I leave it an open question whether this person should be described as someone prone to spells of madness.

of a pair of adjacent school districts one is sent to.) Judgment in particular cases will not always be straightforward. Some will say that the protagonist is lucky enough to know. Others will say the danger of error is close enough as to destroy knowledge in the actual case. But the situation is not crucially different to fake barns cases. There too intuitions will differ in various versions of the case but nonetheless, there will be a range of cases where most of us will be inclined to think that one is not safe enough to know.

B The Danger of Empty Thoughts.

Let us assume, with orthodoxy, that claims containing empty singular terms fail to express propositions. (I set to one side those free logics that validate identity claims even in the absence of a referent.) Suppose I accept the claim 'Muhammed Ali is Muhammed Ali' and thereby judge that Muhammed Ali is Muhammed Ali. By the Independence₁ test, the claim is not a priori, since I have an intrinsic duplicate who does not know (or even think) anything at all by the counterpart episode. It is also easy to contrive an argument that this is not a case of a priori knowledge even by the Independence₂ test. For consider a world where I think the singular thought that Muhammed Ali is Muhammed Ali but where I could very easily have had an empty thought instead. In that case, my singular thought will not be safe (assuming that we should craft the notion of safety so that the risk of a failed attempt at belief is as damning as the risk of false belief).

Now this conclusion may be welcomed by some. Strictly speaking, I do not know a priori that Muhammed Ali is Muhammed Ali. All that I know a priori is that everything is identical with itself. But this conclusion can be welcomed only if it can be contained. For it is at least arguable that if the world is not compliant, very many of our *predicates* will be semantically empty as well. Certainly, the mere fact that an expression is a predicate does not guarantee its having a semantic value. (If 'Arabia' has no semantic value, then surely 'Arabian' will not either.) I lack the space here to explore whether the danger of emptiness is confined to a very limited class of predicates. Let me, for now, offer a conditional conclusion: If the possibility of emptiness is pervasive, then the environment dependence test will, if insisted upon, reckon pretty much all of our knowledge a posteriori.

C: The Social Nature of Meaning

One lesson of semantic externalism is that one's beliefs are constitutively dependent on the practices of one's community. This constitutive dependence (as I am conceiving it) is not an artifact of some deliberate, or even tacit, act of deference on the part of the thinker. Rather it is the inevitable consequence of the thinker's participation in a linguistic community.¹⁴

Once this lesson is absorbed, it is easy enough to make trouble for the environment independence idea. Consider my belief that necessarily, all unmarried male humans are bachelors. Consider a series of cases, each pairwise slightly different, at one

¹⁴ Put picturesquely: a speaker simply lacks the power to lay down a law that his words and thoughts are not going to be constitutively determined by those around him. He simply hasn't got that kind of semantic authority.

end of which is my actual belief, at the other end of which, owing to modal semantic drift induced by modal variance in community practice, the counterpart of my belief that necessarily all unmarried males are bachelors is a false belief. Let us stipulate further that there is no intrinsic variation among my counterparts across the possible situations. We could imagine, for example, that at the actual world there is very little tendency for the population to say that Adam (of the Adam and Eve story) is not a bachelor, but at the other end there is a very strong tendency among the population to make the counterpart of this speech (reasoning that ‘bachelor’ is not true of Adam since there is no institution of marriage).¹⁵ At that world my inclinations with regard to the case are swamped by my community’s in such a way that ‘bachelor’ in my mouth is not true of possible individuals like Adam. At that world, the counterpart of my belief is false and hence, obviously, not knowledge. If I assume the first conception of environment independence then I have to conclude that my actual belief is not a priori knowledge. Meanwhile, we can show that it is not a priori by the second conception by considering a world in the series of cases where I do believe the same as at the actual world but where I am situated so that very small changes in the community would have meant that I believed a false proposition. At that world my belief that necessarily, all unmarried males are bachelors is not safe enough to count as knowledge.

Given a natural version of semantic externalism, there is a further related kind of threat, having to do with diachronic drift. If the contents of my thoughts are constitutively dependent upon community practice, then my thoughts can undergo semantic drift over time without my noticing, due to shifting practices in my community. If standards change without my noticing it, this can in turn be disruptive to my intuitions and reasoning in all sorts of ways. Sentences stored away in lexical memory might at one time express a true proposition and later express a false one. Predicates might shift semantic value during the diachronic process of some long chain of reasoning. And so on. It is in turn not hard to see various ways in which external semantic stability provides a safe haven for various kinds of reasoning that would typically be counted as a priori. Once again we are faced with a dilemma: abandon the environmental independence idea or else admit that hardly any of our knowledge is a priori.

#2 Experience Independence

We have been exploring the idea that a priority is constituted by the irrelevance of the environment to one’s knowledge. The idea has not proven promising. All human knowledge – even knowledge that philosophers naturally classify as a priori – requires that the environment be so configured as to provide a safe haven for the knower. Perhaps, though, one can make something of the idea that knowledge is independent of experience in a way that does not challenge the material just presented. Think of some paradigmatically a priori proof. Grant that the environment plays the epistemological role of providing a safe haven for the knower. Still, it seems that experience of the environment is not an important part of the *process* whereby the knower achieves knowledge of the conclusion of the proof. And that might appear to constitute the basis of a natural division between certain kinds of human knowledge and others, perhaps the

¹⁵ Thanks to Jose Benardete for this example.

most eligible joint in the vicinity of traditional discussions of a priority. Here is a recent and representative proposal in the vicinity:

To say that something can be known without any justification from the character of the subject's experience is to say that there is a way of coming to know it which does not rely on any such justification. ...On this approach, an a priori proposition is one such that there is a way of coming to know it under which the thinker's entitlement to accept the proposition does not involve the character of the experience.¹⁶

Some will find it natural to develop this idea by invoking the notion of *grounds*: A case of knowing is a priori iff experience does not form part of the grounds of the belief. I find the notion of ground a bit unclear, or at least philosophically tendentious. When I form a perceptual belief about the world, it is not obvious that I need do this on the basis of beliefs or knowledge *about* my experience. Relatedly, then, it is not clear that propositions about my experience form part of my evidence. What commitments are taken on board, exactly, when one says that experience is part of the grounds of belief in this case? Consider also a memory based belief about the past configuration of the external world, one that is not accompanied by any present experience. What exactly does it mean to say that experience is the grounds of the belief?

We can bypass the concept of grounds by appealing instead to methods, where a method of belief-formation is a process that delivers or sustains a belief. Now some methods will be experience-involving – they will be processes, one step or element of which will involve perceptual engagement with the outside world. Others will not. This provides us with a method-based conception of a priority.

Experience-Independence: A case of knowing is a priori if it is sustained by a method that is not experience-involving.

On the face of it, this gives us pretty good results. Mathematical knowledge achieved by pure ratiocination counts as a priori, paradigmatically a posteriori knowledge does not. Laxer, derivative notions of a priority can be contrived to handle cases that do not strictly count as a priori by the above test. If I use pen and paper to do long division then, by the current test, my knowledge is not a priori. But the method I use does bear a close resemblance to methods that are not experience involving. In an extended sense – one according to which methods that sufficiently resemble those that are strictly a priori themselves count as a priori – the pen and paper case will count as a priori knowledge.¹⁷

¹⁶ Christopher Peacocke and Paul Boghossian, Introduction to *New Essays on the A Priori*, Clarendon Press, 2000, 1-2.

¹⁷ It is less helpful to say that the knowledge does not *essentially* depend on the pen and paper. If that just means that it is possible to know that proposition without pen and paper, that will not distinguish the pen and paper method from a paradigmatically a posteriori method. If it means that that very method could have occurred without the pen of paper, the thesis takes on tendentious commitments about the individuation of methods that can be bypassed by falling back on the true and less tendentious observation that a relevantly similar method could have occurred with the pen and paper.

It is worth noticing that if we drop the requirement of environment independence and tie a priority to experience independence, then we will thereby make room for the possibility of all sorts of knowledge that is a priori and contingent. Of course, we have been habituated to allow for certain cases of contingent a priori knowledge. If I believe the proposition expressed by ‘Actually there is an odd number of red lego blocks iff there is an odd number of red lego blocks’ on the basis of an appreciation of the semantic workings of ‘Actually’, then it is natural to count that as contingent a priori knowledge.¹⁸ But with the requirement of environment dependence dropped, and the experience independent notion instituted, we will be encouraged to cast the net of contingent a priori much wider.¹⁹ Suppose, for example, that someone is born with an innate mechanism that is a wonderfully reliable source of beliefs about a certain contingent subject matter – phonetics, syntax, the behavior of falling objects, or the psychological patterns of humans. . . . On the safety theoretic conception, such an individual will have not merely innate beliefs, but innate knowledge. And on the experience-dependent conception, such beliefs would count as a priori. That a duplicate might not have knowledge will be no grounds for complaint here – that would be to fall back on the discarded environment dependence idea. And that such an individual would have no *reasons* for his belief would be no grounds for complaint either – on the safety theoretic conception, there is no obstacle to foundational knowledge that has no further propositional basis. However, I don’t take any of this as decisive reason to reject the experience-dependent strategy for discerning a natural a priori/aposteriori divide. The main obstacle to a workable conception of the a priori that is built upon the experience-dependence idea has instead to do with the individuation of methods.

Consider a particular case of believing that p. There are all sorts of channels of cause and effect leading from the near and distant past to that belief. Which are to count as part of the method of forming the belief? Let us take a closer look at some of the pertinent decision points.

(i) It is often noted that even in paradigmatic a priori cases, empirical evidence can destroy knowledge. Even though I have carefully worked through a mathematical proof that p, I will not know p if I get empirical evidence that I am mad, or that human or mechanized experts have agreed that not-p, or that there is a priori gas in the area, or that I have made lots of mistakes using a very similar proof technique in the past or that lots of smart people are inclined to laugh when they hear my proof. Were such experiences part of my history, then certain episodes in which I make a mistake would then (arguably) count as relevantly similar,²⁰ destroying knowledge in the case at hand. Call knowledge-

¹⁸ Of course, as the preceding discussion shows, even this may not count as a priori by the environment dependence test since (i) the relevant semantic workings may not be internally determined and (ii) in any case, one might need a safe haven in order to *know* the semantic workings, a fact that tends to be obscured by the use of verbs like ‘appreciate’.

¹⁹ For relevant discussion, see my ‘Deeply Contingent A Priori Knowledge,’ *Philosophy and Phenomenological Research*, 65, 2002, 247-265.

²⁰ There is more to be said here concerning when evidence that one does not know destroys knowledge. I shall not pursue the topic here. For now, I shall make the simplifying assumption that, in general, powerful evidence that one does not know destroys knowledge. I am not, however, altogether averse to the idea that in a case where one competently proves a result and sticks to one’s guns despite (uncontested) testimonial evidence from acknowledged experts to the contrary, one knows despite having evidence that one does not.

destroying experiences Bad experiences. Call the remainder Good experiences. My stream of experience is not irrelevant to my knowledge that p. That my proof counts as knowledge appears to depend crucially on its being accompanied by Good experiences. But if the process of arriving at putatively a priori knowledge is individuated so as to include Good experiences, then it will count as a posteriori by the experience dependent criterion. (There is a general problem of *width* here. All sorts of facts at one time have some causal bearing on belief at a slightly later time. Which of them are to count as *part of the belief forming process*?)²¹

A natural reaction to this concern is to try to make good on the intuitive distinction between cases in which the presence of a certain kind of experience is epistemologically important versus the cases in which the absence of a certain kind of experience is epistemologically important.²² There are a variety of tricky questions in the vicinity here. Can omissions as well as positive events count as part of a process? If so, should the presence of an experiential omission in a process count as experience-involving in the relevant sense? I shall not pursue these matters further here. For now, we should recognize that without some acceptable resolution of the difficulty, the experience-independent conception will exclude pretty much all of our putatively a priori beliefs.

(ii) Suppose a teacher instructs someone in the laws of nature when he is young and then he remembers those laws later. There is a store in the memory bank that calls forth the laws on demand, applies them to various possible cases, thereby extracting conditional predictions, derived generalizations and so on. (We can think of two versions of the case. On one version, preservative memory of the laws is accompanied by memory of the teacher giving instruction in the laws. On another version, the laws are stored away without accompanying episodic memories.) Let us suppose that the person's nomic beliefs, conditional predictions, and so on, are highly reliable. Just fill out the details of the case in such a way that at least many of the relevant beliefs are safe enough to count as knowledge. (I assume, here of course, that if the mechanism of belief production is safe enough, then the believer will know its deliverances even if he cannot produce any empirical reasons whatsoever for the laws.)

Recall Frege's test for a priority in the Grundlagen:

If on the contrary, its proof can be derived exclusively from general laws, which themselves neither need nor admit proof, then the truth is a priori.²³

If we apply this to the case at hand we will get an affirmative verdict on a priority. Suppose, for example, that the laws are retrieved from memory and applied to deliver a particular conditional prediction. It does not seem that belief in the laws *need* proof to

(It might destroy knowledge that one knows in such a case without destroying the knowledge of the result.) That view would certainly help with the current challenge.

²¹ This obviously relates to the generality problem for reliabilism that is widely discussed in the literature, though there the emphasis is sometimes on the plurality of types that a single process instantiates, rather than the plurality of processes that deliver a belief. (There are various process types instantiated by a process, some more abstract than others, leaving out more or less detail. Construed abstractly, my opening a can with an electric opener will be the same kind of process and my doing so with a manual opener.)

²² Relatedly: Bad experiences lower evidential probability but Good experiences don't raise it.

²³ Gottlob Frege, *The Foundations of Arithmetic*, second revised edition, Northwestern Press, 1980, 4

count as knowledge, at least from the safety-theoretic perspective. Also, it does not seem the laws *admit of* proof from more banal empirical premises. Now of course it would have been utterly anathema to Frege to count the relevant case as one of knowing a priori. And that is because it never occurred to him that one could have foundational knowledge of contingent propositions by virtue of safe mechanism rather than propositional evidence or experience. All of this merely underscores the need to carefully rethink traditional discussions of the apriori in the light of the safety-based picture.

Let us think about the law of nature case afresh. Should we count the believer as knowing the relevant propositions a priori? One will naturally react – on either version of the case – by insisting that the laws were not known a priori. And one might naturally invoke something like the experience dependent conception in justifying such a claim: the process that led to the fixation of belief included experiential exposure to the teacher. The knowledge is a posteriori knowledge, achieved via testimony, not a priori knowledge.

It is important to realize here that when one considers a stream of events, various processes can be distinguished, some longer in temporal extent than others. Suppose a cake is baked. There is a process that begins with shopping for the ingredients and culminates in the taking the cake out of the oven. But there is a shorter process that begins with putting the unbaked cake in the oven and ends with taking it out. So with judgment. Suppose someone extracts a conditional prediction about the course of events. There is a process that begins with the teacher telling him the laws and ends with applying some laws to derive a conditional prediction. But there is a shorter process that begins with retrieving the laws from the relevant internal information bank and ends with producing the conditional prediction. One of the processes is experience dependent. One is not. Which shall we use to test whether the belief is a priori? Let us call the process beginning with the retrieval Short and the process beginning with the interaction with the teacher Long. Is there any deep mistake in taking Short to be the relevant safe method? (We earlier encountered the problem of width; we are now confronted with the problem of *length*.)

It is no use complaining that Short would not be safe were it embedded in a world where the laws were different. That is true of Long as well and, given the failure of environment independence, is in any case not a good test for a priority.

One might argue that it is very artificial to treat Short as the relevant process, since a process like Short brings knowledge with it only in the presence of Long. The form of argument is not immediately persuasive. But in any case, its premise is faulty. Someone who was born with an innate store of the laws of nature would exemplify something like Short in contriving conditional predictions, but not Long. And as we have seen, with environment independence discarded, there is no good reason to disallow some such cases as knowledge. Even in the version of the case where retrieval is accompanied by episodic memory of teaching, it is not clear that Short can deliver knowledge only in the presence of Long. Consider the following thought experiment: Suppose someone is born with an innate storehouse of true mathematical beliefs and reliable computational techniques. Those beliefs are accompanied by pseudo-memories of being taught mathematics in an earlier life. Those pseudo-memories trick the believer into thinking that Plato's doctrine of recollection – or something like it – is true. It seems to me that we may well be inclined to count such a person as knowing mathematics, despite the

misleading pseudo-memories. Suppose now someone is born with a storehouse of beliefs about the laws of nature – accompanied with relevant pseudo-memories of nomic training from a sage. Here we would have Short – in the version from the episodic memory case -- without Long. Once again, it is far from clear that the person fails to know.

(iii) Consider a case of mathematical knowledge. Here again, there are a variety of candidate processes to use as a test for experience independence. Let us focus on two: (a) Math-Short, which begins with the retrieval of mathematical information from preservative memory and ends with the application of computational techniques to answer the problem at hand. (b) Math-Long, which begins with the training in the relevant techniques and provision of the relevant mathematical information at school, home or college, and ends with the application of computational techniques to answer the problem at hand. Using Math-Short as our benchmark, the relevant belief is experience-independent. Using Math-Long as our benchmark, it is not. Now what is it that justifies using Long in the law of nature case but Math-Short in the mathematical case? A certain kind of answer will be natural to many philosophers at this point. They will insist that the training in the mathematical case ‘doesn’t count’ because that training merely played the role of allowing the person in question to grasp the relevant propositions/acquire the relevant concepts. In this connection they will likely invoke some standard platitude about the distinction between conditions of acquisition and conditions of entitlement. Here is a recent affirmation:

In the case of a priori propositions, much experience, perhaps of a specific character, may be required to grasp the concepts implicated in the proposition or to access the entitlement to believe it; but conditions of grasp and of access remain distinct from the nature of entitlement.²⁴

Various concerns might be raised about this reaction to the question under consideration. First, the original training may have gone well beyond that of allowing the person to understand the relevant propositions. If a mathematical belief is reckoned a posteriori in cases where the training goes beyond introduction to the concepts to provide testimony about the objects of the concepts, many of the beliefs in question will turn out a posteriori. Suppose, for example, the training includes various sophisticated proof techniques that are quite obviously not preconditions of understanding anything in the domain. Moreover, once we take the social character of meaning seriously, it is natural to think of the conditions of meaningful participation in the discourse as quite minimal. Someone who was unsure whether ‘1-2’ had an answer or was undefined counts as meaning what everyone else does by ‘1-2’ on account of social facts. (Meanwhile, any suggestion of some notion of ‘real understanding’ according to which such a person means but does not *really grasp* the proposition in question posits a natural kind of real understanding that may very well be chimerical.) The more one thinks in this social externalist way, the more one has to recognize the training in question as going beyond providing the conditions for understanding. One should also wonder whether belief in any of the basic propositions of arithmetic is really a condition for understanding. Would a community of naysayers that were trained to regard arithmetic as false (on the grounds

²⁴ Peacocke and Boghossian, Introduction to *New Essays on the A Priori*, 2.

that Hartry Field is right about numbers²⁵) not count as grasping arithmetical propositions? The attempt to discount the training on the grounds that it merely introduces the person to the relevant concepts is not compelling. And of course it is no good to fall back on the thought that once the training is over one does not need to remember the experience of the training in order to know the math (since one now has one it takes to know the math just by relying on inner resources). For of course the same could be said in the law of nature case.

(iv) As Timothy Williamson has emphasized in some recent work, much of what actually goes on when people describe themselves as having intuitions about a subject matter is that they form judgments about contingent counterfactual propositions.²⁶ Let me briefly summarize a few of the points that he has pressed using a vivid example. Suppose I reflect and judge that were someone to have less than two hairs he would be pretty bald. Now there are worlds where the counterfactual is false: If a person had a single hair that was 50 feet long and wrapped into a bun, he would not be pretty bald, nor would a person with a single hair that was eight inches thick. Now of course we might think that associated with the contingently true counterfactual is some necessarily true proposition. But it is a delicate and challenging task to try to 'clean up' the counterfactual so as to turn it into a necessary truth. What we do in general is make do with the contingent verdict.

Consider someone who makes various contingent counterfactual judgments about bald people, pretty bald people, and very bald people. There is no actual perceptual engagement with the world during the course of the judgment. He merely reflects, as it were, and produces the counterfactual judgment. This pervasive phenomenon forms an excellent kind of test case for the current discussion. If we adopt the experience-dependent test of a priority, our verdict will once again turn on the individuation of methods. To simplify, let us suppose the person has a little module in his head that is responsible for the baldness verdicts, one that has been crafted by exposure to the world, but which is now eminently reliably with regard to such verdicts, and can, indeed be relied upon without resorting to further perceptual consultations with the environment. We can individuate the belief forming process in a short-term way – as involving deployment of this inner module. So construed the belief will be counted as a priori. Or we can individuate it in such a way that the past experiential engagements that were formative of the module were part of the process. So construed we will count the belief as a posteriori. Do we then think there is some deep fact about which process is constitutive of the knowledge? Either way the person counts as safe. So we cannot ground a choice of method on the question whether the person knows the relevant counterfactual. Moreover, it hardly seems that ordinary folk ever think about such questions, even tacitly. It doesn't seem, then, that we then can ground a preference between the short and long versions of the process by appealing to folk practice. Why then be confident at all that there is some natural joint that is being tracked by

²⁵ See *Science without Numbers*, Blackwell, 1980.

²⁶ See his 'Armchair Philosophy, Metaphysical Modality and Counterfactual Thinking,' in *Proceedings of the Aristotelian Society*, 2005, 1-23.

philosophers' talk about 'relying on experience', one that can settle questions of a priority in a (far from unusual) case like this one?²⁷

#3 Justification and the intellect

Some will react to the discussion so far by claiming that the notion of apriori justification is primary. On one natural version of this view, my discussion has not touched on one of the deepest epistemological facts in the vicinity, namely that intellectual intuition is a source of justification. The picture I have in mind reckons there to be some deep analogy between the way that perceptual experience offers justification for perceptual beliefs and the way that intellectual intuition – 'intellectual seemings' as George Bealer puts it²⁸ – justifies various mathematical, logical and metaphysical beliefs.

Begin with the case of perceptual experience. One natural idea is that it is an intrinsic fact about perceptual experience that it lends support to beliefs in the propositions represented by that experience. On such a view, this justificatory role isn't a contingent fact about the experience; it is of the essence of a perceptual experience that represents *p* that it lends support to a belief that *p*. Of course if there is other countervailing evidence, the belief that *p* may, all things considered, not be justified. But *ceteris paribus*, the experience justifies the belief. Put another way, the experience offers *prima facie* justification for the belief. If we use the language of evidential probability to encode how well supported a belief is, we can articulate one consequence of this picture: of necessity, an perceptual experience that *p* raises the evidential probability that *p* unless the evidential probability that *p* is already 1 or unless the experience is accompanied by countervailing evidence.²⁹ Now on the picture that I am currently considering, much of the same structure carries over to intellectual seemings: by their very nature, they offer *prima facie* justification for beliefs and, coordinately, tend to be evidential probability raisers.

Some might see some of the earlier examples through the lens of this kind of picture. Consider a priori gas. One might argue that the risk of exposure defeats knowledge but not the justification delivered by intellectual seemings. (One might go further, arguing that even actual exposure does not destroy justification unless there is evidence that exposure is going on.³⁰) This kind of reaction is most naturally motivated

²⁷ Williamson's own perspective is fully in accord with the line taken here: 'Both crude rationalists and crude empiricists in effect assume that evaluations of counterfactuals can somehow be divided according to their sensitivity or otherwise to experience into two drastically contrasted classes. That assumption is neither intrinsically plausible nor adequately supported by evidence,' 'Armchair Philosophy, Metaphysical Modality and Counterfactual Thinking,' 18.

²⁸ See, for example, 'A Priori Knowledge and the Scope of Philosophy,' *Philosophical Studies*, 1996, 121-142.

²⁹ The picture says more, of course: that in the absence of counterevidence experience raises the evidential probability to a level such that the belief then counts as 'justified' – but I don't wish to focus on that aspect of the view – and its difficulties -- here.

³⁰ One might instead think while intellectual seeming always bring *prima facie* justification, justification is defeated by the mere presence of a priori gas. (In his *In Defense of Pure Reason*, Cambridge University Press, 1998, Bonjour argues for a view of this kind when discussing the justificatory role of apparent rational insight when it occurs in a being whose judgment is 'irreparably clouded' (137)). For my part, I find it quite hard to keep track of the terms of art that figure in such discussions. Rather than ponder directly where the justification in such cases is merely 'prima facie', it seems better to ask what the range

by the idea that a priori knowledge decomposes into an ‘internalist’ component that is accessible to the subject – some kind of intellectual seeming, and which forms the bedrock of justification, and an ‘externalist component’ that includes various reliability conditions

For clarity’s sake, let me address the idea in its most brutish form. (Perhaps there are more nuanced, more compelling versions of the idea. I leave it to others to contrive them.) Suppose there is a propositional attitude – intellectually seeming that p (henceforth Seeming) that manifests itself to inner consciousness by a special kind of intellectual phenomenology (call it Glow). (Let me be neutral on whether Seeming brings Glow with it necessarily or contingently.) The proponent of the view then claims that Seeming is by its nature an evidential probability raiser, at least *ceteris paribus*, and that we are clued into Seemings that occur in ourselves by Glows.³¹ And on the package that I am envisaging, he claims that such Seemings are at least a core component of a priori justification.³²

Here are some salient pressure points (ones that will have already occurred to many readers).

First, what should one say about possible cases (if there are such) where Seeming that P occurs without Glow? Is the evidential probability of P raised in that case, even when the intellectual act is not manifest from within? The picture we are discussing is, quite evidently, driven by an internalist mandate to allow only that which is ‘accessible’ from the inside to count as relevant to justification. To allow Seeming to be a probability raiser in this kind of case is to fail on the mandate, which in turn will call into question the intellectual underpinnings of the picture in the first place. Suppose instead one insists that Seeming can justify only when accompanied by Glow (those who posit a mysterious necessary connection will get this result for free). The shape of the worry does not go away. For consider cases in which one has Glow but are not in a position to know one has Glow.³³ Is Seeming plus Glow still an evidential probability raiser in that case? Why should unnoticed facts about one’s life beneath the skin be able to raise evidential probabilities if unnoticed facts about the outside world cannot?

of the subject’s evidence is in such cases, and in particular whether intellectual seeming still counts as evidence for the proposition presented.

³¹ Of course if we wish to develop a serious theory of evidential probability for mathematical knowledge, we shall have to drop the idealization of logical omniscience common to standard decision theories, as well as (relatedly) the picture of sets of worlds as the objects of knowledge that is also orthodox within decision theory.

³² Some versions of the idea proceed in terms of a more inflated ideology. Laurence Bonjour posits a relation of rational insight wherein one is ‘able to simply see or grasp of apprehend that the proposition is necessary’ (106). If ‘see that’ and ‘apprehend that’ merely mean ‘know that’ then rational insight could not play the desired role of explaining a priori knowledge. Bonjour presumably intends it to be understood as a *sui generis* fundamental, quasi-perceptual, relation holding between a subject and necessary propositions. Supposing there were a fundamental relation of this sort, it could certainly be used to identify a kind of knowledge. But I doubt very much that there is. (I note in passing that Bonjour assumes that a proposition must be necessary in order to be known a priori. I find his condensed discussion and rejection of the contingent a priori on p. 12-3 of *In Defense of Pure Reason* altogether unpersuasive, but here is not the place to pursue the matter.)

³³ For relevant discussion see Williamson, *Knowledge and Its Limits*, chapter 4.

The proponent of the view has a natural retreat: Known Seemings are probability raisers. When one knows that it Seems that p , this tends to raise the probability that p . Henceforth I will assume this version of the view.

Second, it is worth pressing on the modal commitments of the view that known Seemings are, by their very nature, probability raisers. Consider a race of creatures of the following sort. They are so constituted that Seemings come easily. All sorts of wild and wonderful propositions Seem to be the case, and have an accompanying Glow. Ancestors that believed what Seemed to be so didn't last long. Evolution selected for species members that took a more questioning attitude to their Seemings and Glows. (We might, if we wish, imagine that there was a special kind of phenomenology – Superglow (perhaps Glow accompanied by an invigorating tingle) – that was in fact reliable and that the creatures did take Seemings accompanied by Superglow at face value. Imagine some such creature born into the world, and has his very first Glow. He has a Seeming that, say, $5 + 12 = 18$, accompanied by Glow, and takes little doxastic notice of it. (We might imagine that, instinctively, he has the habit of responding to a Seeming that P by filing the question whether P as one to be inquired about later. Perhaps Glows are good indications of questions worth asking for these creatures, even if they are not good indications of the answer.) Its not that the creature knows that Glows are unreliable – he hasn't achieved that kind of self-awareness quite yet. Its just that the creature is so constituted as to be generally unaffected doxastically by Glows. The picture that I have been considering maintains that despite all this Known Glow plus Seeming raises the evidential probability that p , for some p , when the creature Glows for the first time. I do not find this position especially plausible.

Third, one wonders why my current Seemings are given such a privileged role viz a viz a priori justification? Let me note two issues in this connection. First, is there some special epistemic status to my current, here and now Seemings? Suppose I remember a Seeming that p from the past but do not have one at present? Won't that do just as well? Or suppose I become aware that you have a Seeming? Second, and more importantly, what should the proponent of the view think about, say, mathematical knowledge that is unaccompanied by Seemings? After all, it seems hopelessly naïve to suppose that all bona fide mathematical knowledge has to be preceded by the intellectual act of Seeming or the phenomenological fanfare of Glow. I suppose that the proponent of the package I have in mind would think that if there was no evidence from perception (testimony) or memory (memories of testimony or earlier Glows), then these would be cases of knowledge without justification. But this only serves to highlight the burden placed on the notion of 'accessible' in standard internalist ideology. As we noted above, Seemings make their epistemological presence felt by being known. Once known, facts about Seemings justify other beliefs, including ones about the necessary structure of reality (at least according to the picture we are entertaining). But parity of reasoning ought to tell us that mathematical facts make their epistemological presence felt by being known, and that they can then in turn justify other beliefs as well. But once this is conceded, it becomes unclear what *special* role Seemings (as opposed to other known facts) play with regard to justification. If 'accessible' means 'known', the internalist restriction to the 'accessible' as delimiting 'what has to go on' will not in any way force one to typical

internalist conclusions as to range of justificatory resources. But if 'accessible' means something else, I do not know what it is.³⁴

#4 Conclusion

Shifty vocabulary that fails to mark any natural joints may yet prove to be an informative vehicle in context. Suppose there is no natural divide among human beings between the sinners and the saints. Still, in context, someone can convey useful information by describing someone in one of these ways. Similarly, when a philosopher tells me 'I figured it out a priori', that will often provide me with information. In context there will be a range of processes that, for all I know, constitute the etiology of his belief. Some will be comparatively more heavily or more recently experience reliant than others. From his speech, I will be able to rule out certain of the processes as actual. Of course, none of this shows that the category of a priority will prove a fertile one from the point of view of serious epistemological theorizing, or that it marks a natural joint in our epistemological lives. Indeed, its prospects on that score seem rather dim.

³⁴ For relevant discussion of this point, see Williamson's 'On Being Justified in One's Head', forthcoming in a festschrift for Robert Audi.