

Analytic-Synthetic and A Priori-A Posteriori

Brian Weatherson

2016

This article focuses on the distinction between analytic truths and synthetic truths (i.e. every truth that isn't analytic), and between a priori truths and a posteriori truths (i.e. every truth that isn't a priori) in philosophy, beginning with a brief historical survey of work on the two distinctions, their relationship to each other, and to the necessary/contingent distinction. Four important stops in the history are considered: two involving Kant and W. V. O. Quine, and two relating to logical positivism and semantic externalism. The article then examines questions that have been raised about the analytic-synthetic and a priori-a posteriori distinctions, such as whether all distinctively philosophical truths fall on one side of the line and whether the distinction is relevant to philosophy. It also discusses the argument that there is a lot more a priori knowledge than we ever thought, and concludes by describing epistemological accounts of analyticity.

1 History

It's easy to give a rough gloss of the notions of analyticity and a priority.

- Something is an analytic truth iff it is true in virtue of its meaning.
- Something is an a priori truth iff it is knowably true without justification by experience.

And this yields us two distinctions, between analytic truths and synthetic truths (i.e., every truth that isn't analytic), and between a priori truths and a posteriori truths (i.e., every truth that isn't a priori). But fleshing out these distinctions takes some work, as we'll see. Let's start by a quick historical survey of work on the two distinctions, their relationship to each other, and their relationship to the necessary/contingent distinction. There are four important stops in the history.

The distinction between "analytic" and "synthetic" traces back to Kant. He thought that both distinctions in our title (analytic/synthetic and a priori/a posteriori) were real, and that they were not the same distinction. In particular, he held that most interesting philosophical and mathematical claims were synthetic a priori. This is because he (at

least most of the time) worked with a fairly narrow notion of analyticity. A subject-predicate sentence A is B is analytic if “the predicate B belongs to the subject A as something that is (covertly) contained in this concept A ” (Kant 1781/1787/1999, 6). But there can be plenty of a priori truths that do not fall under this narrow category.

Let’s look at one example of current importance. Consider the claim *Whatever is known is true*. It is at least plausible that that is a priori; that we don’t need to look into the world to know that knowledge implies truth. But is it analytic for Kant? It is iff there is an analysis of knowledge, and the analysis is of the form S knows that p iff p is true, and X , for some value of X . Most epistemologists nowadays would reject the idea that there is an analysis of knowledge. But even among those who hold out hope for such an analysis, an argument by Linda Zagzebski (1994) has convinced most people that the analysis cannot be of this form. In particular, Zagzebski argued that if there is such an analysis, X must entail that p is true, eliminating the need for this clause. So *Whatever is known is true* will turn out, by Kantian standards, to be synthetic, even if it is a priori.

Although Kant did not think the two distinctions we are focussing on are equivalent, most scholars take him to have thought all and only necessary truths are a priori knowable. (Though see Strang (2011) for a dissent.) This will be a common theme throughout much of the history.

Our second stop on the history is logical positivism, most clearly represented in English by Ayer (1936). The positivists thought that all three distinctions were in a fairly deep sense equivalent. In particular, they thought all three were very close to the distinction between theorems and non-theorems of logic. The positivists were, self-consciously, building on the tradition of British empiricism. But unlike some empiricists, they didn’t want to insist on an empirical basis for logical and mathematical knowledge.¹ The solution was to build on the logicism about mathematics developed by Frege and Russell.

Borrowing a term from Boghossian (1996), let’s say that a sentence is *Frege-analytic* iff it can be converted to a logical truth by the substitutions of synonyms. The positivists thought that all a priori, necessary and analytic truths were the Frege-analytic truths. Without logicism, this would be wildly implausible, since mathematical truths would be an exception. But Frege and Russell had done enough to make that possibility less worrying.

The positivists’ view has some epistemological attractiveness. How can we know things without having empirical input? And how can we know that some things are true not just in this world, but in all worlds? Well, say the positivists, by knowing the language (which we learn empirically) we learn what sentences are related by the substitution of synonyms. And then the puzzles about knowledge of analytic or necessary truths just reduce to puzzles about the epistemology of logic.

Our third major stop then is with Quine, who questioned both steps of this attempted explanation. First, Quine (1936) noted that the story still needs an epistemol-

¹Actually, the history of what pre-positivist empiricists believed about mathematics is a little more complicated than the standard story. See Whitmore (1945) for some details.

ogy of logic. The obvious expansion of the story told so far won't work. It can't just be by learning the meanings of the logical connectives that we come to learn which the logical truths are. That's because we need to be able to derive the consequences of those meanings, and for that we need logic.

But second, Quine (1951) argued that we have no independent way to make sense of the notion of synonymy that is at the heart of Frege-analyticity. This is the most famous part of Quine's attack on the empiricists epistemology of logic and mathematics, but it isn't the strongest part of it. Indeed, the argument in "Two Dogmas" is both strange and self-undermining.² Quine's primary complaint in that paper about the notions of analyticity, synonymy and meaning is that the only way we have of understanding these notions is in terms of the others. But that would only be a problem if we thought we needed to understand them in terms of something else. Arguably we need not; the notions could be theoretical primitives. And especially if one is a confirmation holist (and part of the point of "Two Dogmas" is to defend holism) we shouldn't worry about circularity cropping up near the core of our epistemology. Relatedly, a naturalist like Quine shouldn't care about whether we can give a definition of terms like 'meaning', but rather about whether it is a useful concept in a science like linguistics or cognitive science.

To avoid attributing an incoherent position to Quine, we should interpret Quine the argument of "Two Dogmas" as part of his larger argument against the appeal to meanings and analyticities.³ Quine's larger point, as defended in (Quine 1960), was that meanings were unnecessary scientific postulates. He thinks that we simply don't need them to explain all the facts about cognition and communication that need explaining. Now it isn't clear how many people will share Quine's view that meanings are unnecessary for these sciences, since without his behaviourism the attempt to do without meanings looks unsuccessful. But the larger point is that Quine isn't simply relying on an argument from the irreducibility of analyticity to a dismissal of the analytic/synthetic distinction.

The last stop on our history tour is semantic externalism. The externalists complicated the above story in two overlapping ways. First, they developed convincing arguments that necessity and a priority were dissociable. The most compelling of these arguments were the examples of necessary a posteriori truths, such as *Water contains oxygen*. No matter what surface characteristics or functional roles a substance might play, if it does not contain oxygen, it could not be water.

Second, they showed that the pre-theoretical notion of meaning, which had seemed good enough for much prior philosophical theorising, contained a number of distinct ideas. Here is how Gillian Russell (whose writings I've leaned heavily on in this introduction) puts it,

In three astonishingly influential pieces of philosophical writing, Hilary Putnam (1973) argued that meaning couldn't be both what a speaker

²The next few sentences follow the arguments of Sober (2000) fairly closely.

³This paragraph follows closely the discussion of Quine in Russell (2008).

grasped and what determined extension, Kaplan (1989) argued that what determines extension (character) and what got contributed to what a sentence said (content) came apart in the cases of indexicals and demonstratives, and Kripke (1980) argued that what determined the extension of a name or natural kind term need not be known in order for a speaker to understand the expression, nor was it what was contributed to the proposition expressed by a sentence containing one. Each was suggesting that the roles attributed to a single thing—the expression’s meaning-in the [pre-theoretical] picture, can be played by distinct things. (Russell 2008 x)

From this point on, when we talk about truth in virtue of meaning, we have to clarify which aspect of ‘meaning’ we mean. With that in mind, let’s turn to the questions that have been raised about the distinctions.

2 Five Questions

To focus our discussion, let’s start with five questions we could ask about either the purported distinction between analytic and synthetic, or between a priori and a posteriori.

1. Is there a sensible distinction here?
2. Are there truths on either side of the line?
3. Does the distinction track something of independent significance?
4. Do all distinctively philosophical truths fall on one side of the line?
5. Is the distinction relevant to philosophy?

The questions are obviously not independent; a negative answer to the first suggests that we better not offer a positive answer to any of the rest, for example. But there are more degrees of freedom here than might immediately be apparent.

A negative answer to the second question, for instance, need not imply a negative answer to the first. If one held, with Phillip Kitcher (1980) that a priori warrant is by its nature indefeasible, and as a matter of fact no warrants are indefeasible, then one would think the a priori/a posteriori distinction is sensible, but in fact everything falls on one side of it.

With respect to the a priori/a posteriori distinction, I’ll argue below that while the answer to question 4 is clearly negative, the answer to question 5 is positive. The a priori/a posteriori distinction may be relevant to philosophy even if it isn’t relevant to, for example, demarcating philosophy from non-philosophy. Alternatively, a positive answer to question 5 may follow from a *negative* answer to one of the earlier questions. (Williamson (2013) suggests, but ultimately I think does not endorse, the view suggested in the following sentences.) If we learned that all knowledge was a posteriori, i.e., that all knowledge depended in an epistemologically significant way on experience, that would be epistemologically interesting. So the distinction might have a valuable

role in articulating and perhaps defending a key philosophical insight, even if all the actual cases fall on one side of the distinction.

The point of raising these questions at the start is to ward off a possible confusion that can easily arise when discussing distinctions. It is common to hear about ‘attacks’ on a distinction, or ‘scepticism’ about a distinction, but a moment’s reflection shows that it isn’t clear what this comes to. I think that most of the ‘attacks’ on either of our two distinctions are arguments for a negative answer to one of these five questions. (We’ll see some instances of this as we go through the entry.) But different attackers may argue for different negative answers, and different defenders defend different positive answers. So it is, I think, helpful to have these distinct questions in mind before we begin.

3 The Traditional Notion of the A Priori

The traditional notion of the a priori makes the best sense, I think, if you start with the following three assumptions.

1. There is a notion of justification that is distinct from, but a constituent of, knowledge.
2. Whether a belief is justified, in this sense, depends just on the evidence the believer has.
3. Evidence about the external world consists solely of perceptual experiences.

From 2 and 3 we get the idea that there could be some beliefs whose justification does not depend on any perceptual experience, i.e., beliefs that are justified by a null set of perceptual experiences. These are the beliefs that are justified first, i.e., a priori. Then by 1 we can say that these beliefs satisfy a part, possibly a large part, of the conditions for being knowledge. And this is the a priori knowledge.

The problem, as will probably be clear to most readers, is that all three of the assumptions I started with are contentious. As noted in the introduction, Linda Zagzebski (1994) has shown that there cannot be any non-factive notion of justification that is a constituent of knowledge. Timothy Williamson (2000, Ch. 8) has argued convincingly against the phenomenal account of evidence; our evidence consists of facts about the world, not just facts about our experience. Point 2 is less clearly mistaken, but is still far from obvious. (See Conee and Feldman (2004) for a long defence of point 2, as well as discussion of several problems with it.)

Once we drop the three ideas though, or even just the first and third, what could be left to say about the a priori? A natural first move is to think about what *explains* a person’s knowledge, rather than what *constitutes* it. On the classical picture I just sketched, Bob’s knowledge that there are tigers nearby might be constituted by his experience of hearing tiger-like growls. On that picture, having that experience is (partially) constitutive of being justified in believing that there are tigers nearby, and that justification is (partially) constitutive of his knowing there are tigers nearby, and it is these constitutive connections that make his knowledge a posteriori. We don’t need to make assumptions

that are nearly so strong to conclude that the experience partially *explains* his knowledge. The experience could (partially) explain why he is justified, without being any part of the justification, and the justification could (partially) explain why he knows, without being any part of the knowledge.

But there's a problem with this move too. I know that all tigers are tigers. On a standard view about the a priori, this will be a piece of a priori knowledge. But to explain why I have that knowledge, you have to appeal to some experiences I have had. After all, with no experiences, I would not be able to think about tigers. So maybe nothing will end up a priori.

There's another usual response here. The experiences I have *enable* me to think about tigers, without doing anything to *justify* my belief that all tigers are tigers. So maybe a priori knowledge is that knowledge where experiences do not play any justificatory role, although they may play an enabling role.

That distinction between justifying and enabling will do a lot of work in what follows, so it is worth pausing over it. Perhaps we can say a bit more precisely what it means. An experience is a mere enabler if it explains why a person knows that *p*, but not in virtue of explaining how it is they can believe that *p*. I think something like that is plausibly true, but it still makes a rather large epistemological assumption, namely that justification is explanatorily prior to knowledge. That's something that will be rejected by those who accept the 'knowledge first' epistemology of Williamson (2000).

If you don't accept that justification is explanatorily prior to knowledge, this route at least to articulating the difference between experiences that enable, and experiences that justify, is closed off. And perhaps the enabling/justifying distinction is too obscure to do much work. That's what Williamson has recently argued, and in the next section we'll look at his argument.

4 A Priori Knowledge and Practical Skills

The a priori/a posteriori distinction, on the best way of freeing it from outdated epistemological assumptions, relies on the idea we can make sense of the idea that some experiences are necessary for knowledge because they *enable* that knowledge, rather than that they *justify* that knowledge. Timothy Williamson (2013) has argued that this distinction is too unclear to be useful, and as a result the a priori/a posteriori distinction cannot do the work epistemologists need.

I find the example Williamson uses, involving Norman and *Who's Who* (Williamson 2013, 295), rather unintuitive, so I'll substitute a different example that I think makes the same point. Diane is a great basketball player. One of her great skills is being able to anticipate the moves a defence will make, and responding with a move that will maximise her team's chance of scoring. This is a skill she's honed through years of practice and competition. And her most common manifestation of it comes in game situations, when she sees an opposing defender and realizes what move will maximise her team's expected points. But she can also manifest this skill 'off-line', when she considers conditional questions of the form *If the opposing team were to do this, what should I do?*

Williamson notes that in some such cases, these questions will be solved through the use of imagination, which is surely right at least for some sense of ‘imagination’. And in these cases, there won’t be any particular experience that justifies the answer. Yet Diane can acquire knowledge by these acts of the imagination. Is the knowledge she gets a priori or a posteriori? Williamson thinks there is no good answer to this question, since Diane’s experiences play a role in honing her skills that goes beyond the enabling role, but this role very different to the role experiences play in classical examples where we can point to a particular experience that justifies the answer.

Now it might seem that there’s a simple move to make here. Diane’s knowledge is obviously a posteriori because it is explained by her years of experience playing basketball. It is a case of (massive) overdetermination, but that doesn’t mean the experiences collectively are not an essential part of the explanation. Williamson’s response is that if we go down this route, some paradigmatic instances of a priori reasoning will turn out to be a posteriori. For instance, our ability to engage in logical reasoning might turn out to be dependent on our ability to track identity of objects (or even just terms) across time. In general, this kind of response threatens to drive the a priori out of philosophy altogether.⁴

This is not, I hasten to note, Williamson’s conclusion. Williamson thinks that some of our logical knowledge is a priori, and Diane’s knowledge is a posteriori, but the salient explanations of how those pieces of knowledge are obtained and sustained are similar in epistemologically salient respects. So he concludes the a priori/a posteriori distinction does not track anything of epistemological significance.

5 Innate Knowledge

In the previous section we considered an argument that there is much less a priori knowledge than we usually assume. In this section we’ll look at an argument, tracing back to work by John Hawthorne (2007) that there is a lot more a priori knowledge than we ever thought, in principle a lot lot more.

Recall that we’ve argued, on pain of losing all a priori knowledge, that we must understand a priori knowledge as knowledge that is in some sense *prior* to experience, not knowledge that is *independent* of experience. So now consider beliefs that really are prior to experience, namely innate beliefs. There is a lot of evidence that neonates have differential reaction to (right way up) human faces than they do to other objects. (See Chien (2011) and Heron-Delaney, Wirth, and Pascalis (2011) for some recent studies on this and citations to many more.) It is natural to explain this by positing an internal representation in the neonate of the structure of human faces; i.e., a belief about how human faces are structured. Since these beliefs are true, and are in a good sense held because they are true, they seem to amount to knowledge. Yet they are clearly not

⁴Note this argument is distinct from the argument Williamson (2007) makes about the role of knowledge of counterfactuals in philosophical reasoning, and its implications for the a priori status of philosophical knowledge. We’ll return to that argument, and the response by Ichikawa and Jarvis (2009) below. The key point is that this argument only turns on the idea that philosophical reasoning might rest on empirically acquired and honed skills.

grounded in, or explained by, the experiences of the neonate. So they look like a priori knowledge.

This is obviously very different to the standard conception of what is a priori knowledge. As we noted at the top, and will expand on in the next section, there is a lot of interest in the possibility of a priori knowledge of contingent truths. But even the most enthusiastic supporters of the a priori don't think we have a priori knowledge of facial structure of conspecifics.

The problem is actually worse than this. We don't normally focus on what is actually known a priori, but what is a priori *knowable*. The reason for this is fairly simple. Most of us cannot know complicated enough multiplications without the aid of empirical evidence.⁵ But this doesn't compromise the idea that mathematical truths are in a deep sense a priori. That's because one could, in principle, know them a priori, even if creatures with small brains or limited skills need assistance from their perceptions.

But if that's right, then we can imagine creatures with all sorts of different innate beliefs. Indeed, for any law about the world, we can imagine a creature who innately believes that law to hold, and whose belief has the right kind of evolutionary explanation for it to count as knowledge. (Why the restriction to laws? Well, beyond that there might be issues about whether the innate beliefs are accidentally true. In any case, I'm not claiming that *only* laws could be known a priori this way.)

There isn't any obvious way out here. I think the best thing to do is to say that when we say that something is a priori knowable, we have to mean that it is a priori knowable for creatures like us. That rules out the possibility of having all the laws be a priori, but at the cost of making some arithmetic truths a posteriori.

The arguments by Williamson and Hawthorne I've discussed in the last two sections challenge the utility of the traditional notion of the a priori. But for the rest of this entry I'll set them aside, and discuss what ways we might modify, or use, the traditional notion should we find responses to these challenges.

6 Substantive A Priori

As we noted in the introduction, a common thread through much of the history of this topic was a belief in a close relationship between a priori and necessity. Most writers take Kant to have treated them as co-extensive notions, the positivists thought they were identical, and Quine took them to suffer from similar defects. It is only with the externalists that we see a gap appearing between the two.

Even once the externalists appear, the gap is not as wide as it may be for two interlocking reasons. The first is that the argument from externalism to the existence of the necessary a posteriori is clearer than the argument from externalism to the existence of

⁵I think that when one carries out multiplication by hand, using the techniques taught at school, the marks on the paper play a justifying and not an enabling role. But arguing for that would be beyond the scope of this entry. It should be less controversial that multiplications carried out by machine give us a posteriori knowledge of the answer.

the contingent a priori. The second is that externalism may only give a 'shallow' distinction between necessity and a priority. Let's take these in turn.

Assuming externalism, we can identify examples of the necessary a posteriori by using familiar natural kind terms. To take a famous example, it is necessary and a posteriori that water is H_2O . It is harder to even identify the contingent a priori. The rough idea is clear enough. We take the characteristics by which ordinary language users identify water, and say it is a priori that water has those characteristics. But what are those characteristics? Water is the stuff which falls from the sky, fills the rivers, lakes and oceans, and so on. Is any one of these a priori? Not really. It could turn out that nothing had all these properties. (Is it really water in the oceans anyway, or is salt water a different substance?) So we could introduce a new term, Chalmers (1996) suggests 'watery', for the long disjunction of conjunctions of properties that a substance must have if we are to identify it as water. Perhaps we come up with a list such that *Water is watery* will be a priori, though since H_2O need not have been watery, it will be contingent. Note we will, at the least, have to introduce new vocabulary to identify this kind of contingent a priority.

The other worry is that the gap opened up here is 'shallow' in the sense of Gareth Evans (1979). Given the way things turned out, water must be H_2O . But in some intuitive sense, things could have turned out differently. (I'm taking the helpful locution 'could have turned out' from Yablo (2002).) It could have turned out that the stuff in the rivers, oceans etc was XYZ. So while it is necessary that water is H_2O , it could have turned out that this not only wasn't necessary, it wasn't even true. If that all sounds plausible to you, you may well think that the a priori truths are all and only those truths which couldn't have turned out to be false. This way of thinking is behind the important *two-dimensionalist* approach. Important works in this tradition, as well as Evans (1979), include Davies and Humberstone (1980), Chalmers (1996) and Jackson (1998).

Now there are significant challenges facing two-dimensionalists, several of which are set out in Block and Stalnaker (1999). But my sense is that several of these challenges are very similar to the challenges facing anyone trying to get an argument from semantic externalism to the contingent a priori. If we could say more clearly what it is for something to be watery, it would be easier to say whether a particular world is one where water turned out to be XYZ. So I suspect if externalism gives us a reason to believe in the contingent a priori, it will be a fairly shallow distinction. (This doesn't extend to the argument from externalism to the necessary a posteriori; we don't need to shore up two-dimensionalism to say that *Water contains oxygen* is necessary a posteriori.)

But that's not the only way that a priority might outrun necessity. In recent years there has been a surge of interest in the idea that we can know a priori various anti-sceptical propositions. This idea was advanced in detail by John Hawthorne (2002), and then suggested as a way out of sceptical problems by Roger White (2006) and Brian Weatherson (2005).

Recently, Stewart Cohen (2010) and Sinan Dogramaci (2010) have suggested that (assuming inductive scepticism is false), we can use ampliative inferential steps in sup-

positional reasoning. That is, if it possible to inductively infer B when we know A , it is possible to infer B on the supposition that A , and go on to infer the material conditional $A \supset B$. That conditional will be contingent if the inference was ampliative, but since we've discharged the only supposition we used, it could be a priori in a good sense. I have doubts about this route to the contingent a priori (Weatherson 2012), but I think the general idea is plausible.

To make things more concrete, consider 'bubble worlds'. A bubble world consists of a person, and the space immediately around them. If you think that evidence supervenes on sensory irritation, then you have a duplicate in a bubble world who has the same evidence as you.⁶ But you're not in a bubble world, and you know it. There's a well known probabilistic argument that your evidence can't be grounds for ruling out possibilities that entail you have just that evidence.⁷ So your evidence doesn't rule out that you're in a bubble world. But you know you're not. Hence that knowledge is a priori. So *I'm not in a bubble world* might be contingent a priori.⁸ Moreover, it's not a 'shallow' contingency. It could have turned out that you were in a bubble world. Indeed, with some more evidence you might even know this.

I'm not going to defend here the claim that it's contingent a priori that you're not in a bubble world. Indeed, I don't even believe the probabilistic argument for that conclusion that I just referenced. But it is worth noting this trend towards taking seriously the possibility of substantive a priori knowledge.

7 The A Priori In Philosophy

I've argued so far that the best sense we can make of the a priori allows for a lot of a priori knowledge. Once we realise that a priori knowledge, like any other kind of knowledge, is defeasible, and fallible, it seems possible that an agent could have a lot of foundational knowledge of contingent matters. And that foundational knowledge does seem to be a priori. Of course, such an agent would not be very much like *us*; so there is still a question of what agents like us could know a priori. And it might seem that the class of such pieces of knowledge might be relatively small and interesting.

In particular, one might think that philosophical knowledge, or at least some interesting part of philosophical knowledge, might be a priori. Herman Cappelen (2012) notes that a wide range of philosophers, with very different commitments, end up with the view that the a priori has a distinctive role to play in philosophy. (See, especially, chapters 1 and 6 of that book.) But, as Cappelen also shows, these philosophers are

⁶If you prefer a wider conception of evidence, so that for instance two people who are looking at distinct duplicates have distinct evidence, just make the bubble a little bigger, and this argument will still go through.

⁷See White (2006). David Jehle and I have argued that this argument uses distinctively classical logical principles in a way that might be problematic (Jehle and Weatherson 2012). And I've argued that even slight weakenings of the assumptions about how to update credences make the argument fail. (Weatherson 2007).

⁸I'm assuming here that the semantic response to scepticism, as defended by (Hillary Putnam 1981), doesn't work for ruling out bubble worlds. Defending this assumption would take us too far from the current topic.

mistaken; outside perhaps of philosophical logic, the a priori doesn't have a particularly special role to play in philosophical inquiry. We can see this, I think, by working through one recent debate.

Timothy Williamson (2007) noted that philosophical thought experiments are almost always incomplete. The text of an example doesn't guarantee that conclusions that are usually drawn from it. To use his example, to guarantee that the subject in one of Gettier's examples has justified beliefs (that don't amount to knowledge), we have to suppose that there are no defeaters in the vicinity, but that isn't stated in the example. Williamson's solution to this is that we should read the example as a certain kind of counterfactual. What we know, Williamson argues, is that in the nearest world where the example was instantiated, the subject would have justified true belief without knowledge. (I'm ignoring here some complications involving names, and donkey anaphora, that are not relevant to this debate.)

Jonathan Ichikawa and Benjamin Jarvis (2009) object that this makes philosophical knowledge a posteriori. We have to know what the world is like to know what the nearest world in which the Gettier case is instantiated is like. Ichikawa and Jarvis reject this because they want to defend a "traditional" conception of thought experiments on which they provide a priori knowledge. I'm not convinced that this really is part of philosophical tradition; it seems to me the thought experiments in Hobbes, Hume, Mill and many others in the canon rely on empirical knowledge. But I won't press that point here. If one does want to avoid empirical knowledge coming in via the route Williamson suggests, Ichikawa and Jarvis develop a nice way of doing so.

They say that thought experiments are little fictions. We need some empirical knowledge to interpret the fiction. But, they insist, once we are given the fiction, it is a priori that the fiction is possible, and that in it the subject has a justified true belief without knowledge. And, note this point, it is a priori that these two facts entail that it is not necessary that all justified true beliefs amount to knowledge. On the last point, they agree with Williamson. It is also a priori that if the Gettier case could happen, and if it were to happen then there would be justified true belief without knowledge, then it is not necessary that all justified true beliefs are knowledge.

But let's try and generalise this to other thought experiments. Start, for example, with the famous violinist described by Thomson (1971). That violinist plays a key role in an argument whose conclusion is that abortion is often morally permissible. And one premise of the argument is something about an imagined violinist. Williamson will say that premise is an a posteriori counterfactual proposition. Ichikawa and Jarvis will say it is an a priori proposition about what's true in a fiction. Perhaps there's another premise about the possibility of the example, and maybe that's a priori too. But those premises don't come close to supporting Thomson's conclusion. We need another premise about the analogy between the violinist and a woman contemplating having an abortion to get Thomson's conclusion. Any such premise will not be a priori, unless some detailed facts about human biology are a priori. It's a little tricky, but it isn't obvious the soundness of the abductive inference from those premises to Thomson's conclusion is a priori either. (See Pargetter and Bigelow (1997) for some discussion of

this point.)

I think Thomson's example is more typical of philosophical reasoning than Gettier's. We don't just use thought experiments to dismiss theories, like the JTB theory of knowledge. We also use them to defend philosophical conclusions, such as the permissibility of abortion. And in general inferences from a thought experiment to the truth of a theory will involve some a posteriori steps. So even if Ichikawa and Jarvis are right that we can know a lot about thought experiments a priori, it won't follow that in general philosophical knowledge derived from thought experiments is a priori. Get away from special cases where the facts about the thought experiment entail the philosophically interesting result, and this should be reasonably clear.

That doesn't mean that there's no use for the a priori in philosophy. It might be a very helpful concept to use in argument, even if it isn't true that our conclusions are generally a priori. I'll illustrate with one example from my own work. One way to support the sceptical intuition that we don't know we aren't brains in vats is to ask, how could one possibly know that? Rhetorical questions are not arguments, the received wisdom of undergraduates notwithstanding, so the sceptic needs to find some way to extract argumentative force from that question. An attractive option is argue that *these* are all the ways to know something, and you can't know you're not a brain in a vat any of *these* ways. Such an argument typically runs into problems at the first step; arguing that one has exhausted all possible ways of getting knowledge is not easy.

Hume (1739/1978) had the best idea for how to overcome this step. Don't list the ways someone can know something; use some property of knowledge gathering methods to partition the methods. Then argue that in no cell of the partition can one find a method that allows knowledge of the undesired kind. If the partition just consists of the presence or absence of some property, you're guaranteed at least to have covered the field. I've argued Weatherston (2007) that you get an interesting argument by letting the property in question be *is an a priori method*. By 'interesting' I certainly don't mean sound. (And nor do I insist that Hume equated interesting sceptical arguments with sound ones.) But I think you get epistemological insight by thinking about whether knowledge that we're not brains in vats could be a priori, or could be a posteriori. You don't have to think that philosophical conclusions themselves are a priori to think this could be a useful philosophical approach. That last point is probably obvious. We could have developed the sceptical argument by asking whether knowledge of non-vattedness is innate or acquired. But suggesting that's an interesting argument wouldn't imply the very traditional view that philosophical knowledge is typically innate.

8 Metaphysical Accounts of Analyticity

Let's turn now from the a priori to the analytic. As we noted at the start, the traditional notion is that a sentence is analytic iff it is true in virtue of meaning. And, as we saw at the end of the introduction, this notion is complicated by the fact that traditional theories of meaning conflated several things that should be kept separate.

Paul Boghossian (1996) makes a distinction that has been highly influential between

metaphysical and **epistemological** conceptions of analyticity, and I will follow many contemporary writers in splitting the topic up in this way. The metaphysical conception is the one most continuous with the traditional notion of analyticity, and also the one least popular with contemporary theorists, so we will start with that. It is the notion that some sentences are true merely in virtue of their meaning.

Boghossian, following Quine, argues that this notion is either nonsensical or trivial. Consider a simple example of a putatively analytic truth, say *Everything is self-identical*. Why is this true? In part, because it means that everything is self-identical. But that can't be what we mean to say that it is analytic. *Paris is beautiful* is true in part because it means that Paris is beautiful, but that doesn't make the sentence an analytic truth. What we need is that this is the only thing needed for the sentence to be true. And that isn't the case for either sentence. *Everything is self-identical* is true because of what it means and the fact that everything is indeed self-identical, and *Paris is beautiful* is true because of what it means and the fact that Paris is indeed beautiful. We haven't yet found a difference between the two.

It might be easy to see a response here. Start with a less discriminating treatment of truth makers than I supposed in the previous paragraph. Say that a sentence is true in virtue of what it means, and the way the world is. So both of our examples are true in virtue of their meaning and the way this world is. But for *Everything is self-identical*, it doesn't matter how the world actually is, any way it could have been would have made the sentence true. The contribution of the world is like the contribution of the 5 in *What is 0 times 5?* You need a second number there, or the question doesn't make sense, but it doesn't matter which. In some good sense, the 0 does all the work. (This example, and most of the discussion in the rest of this section, lean heavily on chapter 2 of (Russell 2008).)

But this won't do as a conception of analyticity either, because of the examples of the necessary a posteriori. Consider the example *Gold has atomic number 79*. It is true in virtue of what it means, that gold has atomic number 79, and how the world is. But it doesn't matter which world we choose; in any world gold has atomic number 79. Yet it is not, intuitively, analytic.

Russell suggests a solution to this problem that draws on the developments in externalist theories of meaning that we discussed in the introduction. Start with the following three way distinction. (These definitions are a quote from page x of (Russell 2008).)

- **Character:** The thing speakers must know (perhaps tacitly) to count as understanding an expression.
- **Content:** What the word contributes to what a sentence containing it says (the proposition it expresses).
- **Reference Determiner:** A condition which an object must meet in order to be the reference of, or fall in the extension of, an expression.

These can all come apart. In the case of pure indexicals like *I*, the content comes apart from the character and reference determiner in familiar ways. But it is tempting in those

cases to equate character and reference determiner. What makes it the case that a token of *I* picks out me is that I use it, and that relation between usage and content is what someone must know to understand the term. But that's an all too special case. I can be a competent user of the name 'Alex' as a name for my friend Alex without knowing whether she got that name at birth in the normal way, or knowing whether she acquired it later. Competence may require that I know the reference of Alex was somehow determined to be her, but I need not know what that reference determiner was.

Moreover, the character and reference determiner relate to contexts in different ways. It is a familiar point that a sentence like *If you were speaking, I would have been speaking* may be false. That's because when we evaluate the *I* in the consequent, we don't look to who the speaker is in the context of evaluation, i.e., the world where you are speaking, but to the world of utterance, i.e., the context of my utterance. Just like this familiar distinction between contexts of utterance and contexts of evaluation, Russell requires us to think about contexts of introduction. An example helps bring this out.

Say I, on Monday, introduce the name 'Inigo' for the shortest sword fighter. When the name is used on Tuesday, it need not pick out the shortest sword fighter even in the context of utterance. Inigo might have grown, or a shorter person may have taken up sword fighting. Of course, when I use the name in counterfactuals, it might pick out someone who was never a sword fighter. So we distinguish the context the term was introduced in, in this case Monday, from the context it is uttered in, in this case Tuesday.

With these distinctions in mind, we can give Russell's first pass at a definition of analyticity.

A sentence *S* is true in virtue of meaning just in case for all pairs of context of introduction and context of utterance, the proposition expressed by *S* with respect to those contexts is true in the context of evaluation. (Russell 2008, 56)

This will, says Russell, solve the problem about gold. It is true that when someone now utters *Gold has atomic number 79*, they express a necessary truth. But we could have introduced the terms in the very same way, and had the world not cooperated, this sentence would have been false. Indeed, Russell splits analyticity from necessity twice over. She thinks that *I am here now* will be analytic in this sense though it expresses a contingent proposition.

This does rely on understanding what it is for terms in different worlds to have the same reference determiner. Perhaps one could object that had we been pointing at something else when we introduced the term *gold*, we would have been using a crucially different reference determiner. But the issues here about the metaphysics of words and demonstrations, are subtle, and Russell's view that the same reference determiner could determine different contents in different worlds seems plausible.

Russell says that this is a perfectly good notion of truth in virtue of meaning. Of course, as she says, it is really a kind of truth in virtue of what determines meaning, not meaning itself. Reference determiners are part of meta-semantics, not semantics. But

that seems continuous enough with the tradition. And there is no reason to think that analytic sentences, so understood, will be epistemologically distinctive. In this respect we may end up agreeing with the primary conclusion of the discussion of metaphysical analyticity in chapter 3 of Williamson (2007), namely that it isn't directly relevant to philosophical methodology. But it could be an interesting notion in its own right, and as discussed in the previous section, it could be philosophically useful without playing its traditional role.

I have considerably simplified the presentation of Russell's view, however. The definition so far implies that some theorems of geometry, and perhaps fundamental laws of ethics, will be analytic. Like Kant, Russell wants these to be synthetic. Her solution is to say that analytic truths must not just satisfy the constraint given above, but that they must do so because the reference determiners of their parts stand in the right kind of containment relations. But spelling this part of her view out will take too much space, so instead I'll close with a discussion of epistemological analyticity.

9 Epistemological Accounts of Analyticity

In a series of influential articles, Paul (Boghossian 1996, 1997, 2003) argued that we should accept that Quine's argument against metaphysical versions of the analytic/synthetic distinction, but that Quine's arguments left untouched an **epistemic** understanding of the distinction. On this way of understanding the distinction, a sentence is analytic iff it is knowably true merely in virtue of understanding it. Consider, for instance, this sentence.

(E) If frogs bark and ducks howl, then frogs bark.

Now make the following four assumptions.

1. Understanding the non-logical terms in (E) suffices to see it is of the form $(A \wedge B) \rightarrow A$.
2. For logical terms like \wedge and \rightarrow , understanding involves accepting, perhaps implicitly in one's inferential practices, the basic introduction and elimination rules they license.
3. For \wedge , the basic rules are the familiar introduction and elimination rules.
4. For \rightarrow , the basic rules are modus ponens and conditional proof.

Then anyone who understands (E) is in a position to prove it to be true by a trivial three line proof. Generalizing this example, we can get an argument that all logically true sentences are analytic. Generalizing a bit further, we may be able to argue that the propositions they express are a priori knowable, but this requires resolving many of the issues we have already discussed in the discussion of the a priori, and I will set it aside for the remainder of this entry.⁹

⁹There is an interesting worry around here that the three-line argument for (E) is circular, and so cannot justify (E), and this fact undermines Boghossian's argument that it is a priori. See Ebert (2005) and Jenkins (2008) for two ways of developing this worry.

The problem we will focus on is that assumptions 2 and 4, and hence presumably 3, are not clearly true. Actually 4 as stated is almost surely false. If the basic rules for \rightarrow are modus ponens and conditional proof, then \rightarrow is material implication. But \rightarrow was meant to be our symbol for natural language ‘if’, which is not material implication. So the rule must be something else. It isn’t clear what this rule could be. It is plausible that we can use a restricted version of conditional proof when reasoning about ‘if’, such as a version which requires that there be no undischarged assumptions when we apply conditional proof. That will make the proof of (E) go through, but it is unlikely to be a basic rule in the relevant sense, since it does not combine with an elimination rule (i.e., modus ponens) to pick out a unique meaning for ‘if’.

Disagreement about the introduction rule for ‘if’ is endemic to the literature on conditionals. But there is almost a consensus that the elimination rule is modus ponens. Almost, but not quite - Vann McGee (1985) is a notable dissenter. Timothy Williamson (2007) uses the existence of notable dissenters like McGee to mount a sustained assault on Boghossian’s position. It is a consequence of the assumptions we have made, and which Boghossian needs, that anyone who doesn’t accept modus ponens does not understand ‘if’. But that seems implausible. By any familiar standard, McGee understands conditionals quite well. Indeed, he is an expert on them.

This point generalizes, as Williamson stresses. On the inferentialist view about the meaning of logical terms, in any debate about the correctness of fundamental logical principles, either one party doesn’t understand the key terms, or the parties are speaking at cross purposes. The intuitionist mathematician endorses the sentence “All functions are continuous”, and the classical mathematician rejects it. But it isn’t plausible that one party fails to understand ‘all’, ‘functions’ or ‘continuous’, or that they are speaking at cross purposes in that they are assigning different meanings to one of these terms. (I’m assuming the context makes it clear that both parties are speaking of functions whose domain is the reals, and whose range is subset of the reals.) I’ve used an example from real analysis, but we could make the same point less pithily using Pierce’s Law if we wanted to stick to propositional logic.

I’ll close with two replies on behalf of the defender of epistemic analyticity, and some reasons for being dissatisfied with each. The discussion will follow somewhat the recent exchange between Boghossian (2011) and Williamson (2011).

The first response says that we shouldn’t have said if a sentence is epistemically analytic, then *understanding* it is sufficient for knowing that a sentence is true. Rather, we should have said that *knowing the meaning* is sufficient for knowing the sentence is true. Notably, Boghossian does not make this defence in his exchange with Williamson, so it seems he accepts that Williamson was right to take epistemic analyticity to involve a connection between understanding and knowability. And this seems to be right. Consider again *Water contains oxygen*. In one sense of meaning, the meaning of ‘water’ is H_2O . So anyone who knows what ‘water’ means in that sense knows that *Water contains oxygen* is true. But it doesn’t feel like this claim is analytic, especially not in the epistemic sense that interests Boghossian. It is possible that there is some other sense of meaning that will be more useful for Boghossian’s project, but it isn’t clear that know-

ing the meaning in this other sense will differ particularly from understanding.

The second response, and one that Boghossian has used on several occasions, is that the only plausible theory of meaning for the logical connectives is inferentialist, and on an inferentialist theory of meaning it will be true that anyone who understands a connective is disposed to reason correctly with it. That last sentence is deliberately sloppy, much more so than any statement of the response in Boghossian's own work. But the sloppiness is there because it makes a potential equivocation more easily visible.

Consider a theory that says the meaning of a logical connective is either constituted by, or at least constitutively connected to, its appropriate inferential rules. But to understand the term is not to grasp the meaning in this sense, any more than to understand the term 'water' one has to know it is H₂O. Rather, understanding involves participating in the right kind of way in a social practice, and it is that social practice (plus perhaps some facts about the nature of logic, if such facts there be) that determines the appropriate inferential rules for the connective.

Is the theory in the previous paragraph inferentialist? If not, then it is false that no theory other than inferentialism is plausible as an account of the meaning of the logical connectives. For this kind of socialised theory of meaning is, it seems to me, highly plausible. (Williamson (2011) notes that a socialised theory of meaning for the connectives is plausible, though I don't think he would sign up for the view that the result of such socialization is a theory in terms of inferential rules.) If, on the other hand, the theory is inferentialist, then it doesn't follow that understanding requires a disposition to use the rules. Perhaps understanding requires being part of a community many members of which have the appropriate dispositions, but it does not require that any one member have these dispositions. So it won't be true that mere understanding puts one in a position to know. At best, understanding a logical truth means one is in a community in which some people are in a position to the sentence is true. But that doesn't do much to rescue the notion of epistemic analyticity.

References

- Ayer, Alfred. 1936. *Language, Truth and Logic*. London: Gollantz.
- Block, Ned, and Robert Stalnaker. 1999. "Conceptual Analysis, Dualism, and the Explanatory Gap." *Philosophical Review* 108 (1): 1–46. doi: 10.2307/2998259.
- Boghossian, Paul. 1996. "Analyticity Reconsidered." *Noûs* 30 (3): 360–91. doi: 10.2307/2216275.
- . 1997. "Analyticity." In *A Companion to the Philosophy of Language*, edited by Bob Hale and Crispin Wright, 331–68. Oxford: Blackwell.
- . 2003. "Epistemic Analyticity: A Defense." *Grazer Philosophische Studien* 66 (1): 15–35. doi: 10.1163/18756735-90000810.
- . 2011. "Williamson on the a Priori and the Analytic." *Philosophy and Phenomenological Research* 82 (2): 488–97. doi: 10.1111/j.1933-1592.2010.00395.x.
- Cappelen, Herman. 2012. *Philosophy Without Intuitions*. Oxford: Oxford University Press.

- Chalmers, David. 1996. *The Conscious Mind*. Oxford: Oxford University Press.
- Chien, Sarina Hui-Lin. 2011. "No More Top-Heavy Bias: Infants and Adults Prefer Upright Faces but Not Top-Heavy Geometric or Face-Like Patterns." *Journal of Vision* 11 (13): 1–14. doi: 10.1167/11.6.13.
- Cohen, Stewart. 2010. "Bootstrapping, Defeasible Reasoning and *a Priori* Justification." *Philosophical Perspectives* 24 (1): 141–59. doi: 10.1111/j.1520-8583.2010.00188.x.
- Conee, Earl, and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.
- Davies, Martin, and I. L. Humberstone. 1980. "Two Notions of Necessity." *Philosophical Studies* 38 (1): 1–30. doi: 10.1007/bf00354523.
- Dogramaci, Sinan. 2010. "Knowledge of Validity." *Noûs* 44 (3): 403–32. doi: 10.1111/j.1468-0068.2010.00746.x.
- Ebert, Philip. 2005. "Transmission of Warrant Failure and the Notion of Epistemic Analyticity." *Australasian Journal of Philosophy* 83 (4): 505–22. doi: 10.1080/00048400500338724.
- Evans, Gareth. 1979. "Reference and Contingency." *Monist* 62: 161–89.
- Hawthorne, John. 2002. "Deeply Contingent *a Priori* Knowledge." *Philosophy and Phenomenological Research* 65 (2): 247–69. doi: 10.1111/j.1933-1592.2002.tb00201.x.
- . 2007. "Craziness and Metasemantics." *Philosophical Review* 116 (3): 427–40. doi: 10.1215/00318108-2007-004.
- Heron-Delaney, Michelle, Sylvia Wirth, and Olivier Pascalis. 2011. "Infants' Knowledge of Their Own Species." *Philosophical Transactions of the Royal Society B* 366 (1571): 1753–63. doi: 10.1098/rstb.2010.0371.
- Hume, David. 1739/1978. *A Treatise on Human Nature*. Edited by L. A. Selby-Bigge and P. H. Nidditch. Second. Oxford: Clarendon Press.
- Ichikawa, Jonathan, and Benjamin Jarvis. 2009. "Thought-Experiment Intuitions and Truth in Fiction." *Philosophical Studies* 142 (2): 221–46. doi: 10.1007/s11098-007-9184-y.
- Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Clarendon Press: Oxford.
- Jehle, David, and Brian Weatherson. 2012. "Dogmatism, Probability and Logical Uncertainty." In *New Waves in Philosophical Logic*, edited by Greg Restall and Gillian Russell, 95–111. Basingstoke: Palgrave Macmillan.
- Jenkins, C. S. 2008. "Boghossian and Epistemic Analyticity." *Croatian Journal of Philosophy* 8 (22): 113–27.
- Kant, Immanuel. 1781/1787/1999. *Critique of Pure Reason*. Edited by Paul Guyer and Allen Wood. Cambridge: Cambridge University Press.
- Kaplan, David. 1989. "Demonstratives." In *Themes from Kaplan*, edited by Joseph Almog, John Perry, and Howard Wettstein, 481–563. Oxford: Oxford University Press.
- Kitcher, Philip. 1980. "*A Priori* Knowledge." *Philosophical Review* 89 (1): 3–23. doi:

10.2307/2184861.

- Kripke, Saul. 1980. *Naming and Necessity*. Cambridge: Harvard University Press.
- McGee, Vann. 1985. "A Counterexample to Modus Ponens." *Journal of Philosophy* 82 (9): 462–71. doi: 10.2307/2026276.
- Pargetter, Robert, and John Bigelow. 1997. "The Validation of Induction." *Australasian Journal of Philosophy* 75 (1): 62–76. doi: 10.1080/00048409712347671.
- Putnam, Hilary. 1973. "Meaning and Reference." *Journal of Philosophy* 70 (19): 699–711. doi: 10.2307/2025079.
- Putnam, Hillary. 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Quine, W. V. O. 1936. "Truth by Convention." In *Philosophical Essays for a. N. Whitehead*, edited by O. H. Lee, 90–124. New York: Longmans.
- . 1951. "Two Dogmas of Empiricism." *Philosophical Review* 60 (1): 20–43. doi: 10.2307/2181906.
- . 1960. *Word and Object*. Cambridge, MA.: MIT Press.
- Russell, Gillian. 2008. *Truth in Virtue of Meaning: A Defence of the Analytic/Synthetic Distinction*. Oxford: Oxford University Press.
- Sober, Elliot. 2000. "Quine's Two Dogmas." *Aristotelian Society Supplementary Volume* 74 (1): 237–80. doi: 10.1111/1467-8349.00071.
- Strang, Nicholas F. 2011. "Did Kant Conflate the Necessary and the *a Priori*." *Notis* 45 (3): 443–71. doi: 10.1111/j.1468-0068.2010.00809.x.
- Thomson, Judith Jarvis. 1971. "A Defense of Abortion." *Philosophy and Public Affairs* 1 (1): 47–66.
- Weatherson, Brian. 2005. "Scepticism, Rationalism and Externalism." *Oxford Studies in Epistemology* 1: 311–31.
- . 2007. "The Bayesian and the Dogmatist." *Proceedings of the Aristotelian Society* 107: 169–85. doi: 10.1111/j.1467-9264.2007.00217.x.
- . 2012. "Induction and Supposition." *The Reasoner* 6 (6): 78–80.
- White, Roger. 2006. "Problems for Dogmatism." *Philosophical Studies* 131 (3): 525–57. doi: 10.1007/s11098-004-7487-9.
- Whitmore, Charles E. 1945. "Mill and Mathematics: An Historical Note." *Journal of the History of Ideas* 6 (1): 109–12. doi: 10.2307/2707061.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.
- . 2007. *The Philosophy of Philosophy*. Blackwell.
- . 2011. "Reply to Boghossian." *Philosophy and Phenomenological Research* 82 (2): 498–506. doi: 10.1111/j.1933-1592.2010.00400.x.
- . 2013. "How Deep Is the Distinction Between a Priori and a Posteriori Knowledge." In *The a Priori in Philosophy*, edited by Albert Casullo and Joshua C. Thurow, 291–312. Oxford: Oxford University Press.
- Yablo, Stephen. 2002. "Coulda, Woulda, Shoulda." In *Conceivability and Possibility*, edited by Tamar Szabó Gendler and John Hawthorne, 441–92. Oxford: Oxford University Press.
- Zagzebski, Linda. 1994. "The Inescapability of Gettier Problems." *The Philosophical*

Quarterly 44 (174): 65–73. doi: 10.2307/2220147.

Published in *Oxford Handbook of Philosophical Methodology*, 2016, pp. 231-248.