The Distinction works well enough in an elementary introduction to epistemology, for a first-pass survey of the material to be understood. If we wish to go much deeper, the Distinction becomes an impediment to understanding.[3]

B. The Significance of A Priori Justification

Paul Boghossian

B.1 Skepticism About the Distinction

Philosophy has long given pride of place to the distinction between the a priori and the a posteriori (henceforth, "the Distinction"). It has been central not only to its account of many important domains – including logic, mathematics, rationality and morality – but to its very own self-conception. Discarding or demoting the distinction, then, would affect philosophy in multiple profound ways.

Of course, having a venerable history doesn't *guarantee* that a distinction is either well-defined, real, or important. It could turn out that, once the notions ingredient in the Distinction are fully clarified and elaborated, there is not a clear enough conceptual distinction between the a priori and the a posteriori; or that there are no instances of one of its relata; or that, while well-defined and exemplified (real), the distinction is unimportant for epistemology. All of these are possible threats.

What the venerable history does do, though, is create a presumption in favor of the distinction, especially when it is backed up, as it is in this case, by a remarkable agreement about how the Distinction is to be applied. Once it is explained that what we are getting at is a distinction between two ways in which a proposition might be justifiably believed[4] – depending, to put it crudely for now, on whether experience is needed to explain that justification – we seem to have no difficulty in noncollusively agreeing on which beliefs fall on which side of that rough and ready distinction. For example, we are all agreed that the beliefs that

(i) If *p*, then *p*.
(ii) All squares have four sides.
(iii) Nothing can be red and green all over at the same time.
(iv) In a Trolley Scenario, it is morally wrong to throw the heavy man off the bridge in order to save the five innocents trapped below.
(v) Two circles of different radii can at most intersect at two points.
(vi) "I am here now" (as said by someone on a given occasion) expresses a truth.

(typically) all fall on the a priori side,[5] while

[3] I was first alerted to problems with the Distinction through reflection on the epistemology of counterfactuals (Williamson 2007: 165–169). I developed the argument further in Williamson (2013). See Boghossian and Williamson (2020) for an extended debate, and Casullo (forthcoming), Williamson (forthcoming a,b) and Melis and Wright (2022) for further exchanges. Many thanks to Joel David Hamkins, Daniel Kodsi, Jennifer Nagel, and audiences at the universities of Edinburgh, Lisbon, Rijeka, St. Andrews, Texas Tech, and Turin (where it was given as the 2021 LLC Lecture) and the inaugural meeting of the British Society for the Theory of Knowledge at the 2022 Pacific Division meeting of the APA for useful questions and comments on earlier versions.

[4] I will focus on justification rather than knowledge since I regard justification as the more fundamental notion.

[5] The parenthetical "typically" is there to account for the fact that most any proposition that can be learned a priori can also be learned by testimony, which would provide it with an a posteriori justification.

UNCORRECTED PROOFS

c08.indd 183 10-30-2023 17:43:17

(vii) All wooden tables are squares.
(viii) This volume of *Who's Who* is red.
(ix) In all societies, it is considered morally wrong to throw the heavy man off the bridge in a Trolley Scenario.
(x) There are two intersecting circles on this page.
(xi) Whoever knows *p* enjoys knowing *p*.

All fall on the a posteriori side. And we have no trouble knowing which further beliefs to add to which list, no matter how diverse and novel they may be.

Naturally, if the Distinction is to cement its place as a central pillar of epistemology, it has to be formulated more precisely than we have done so far, something I will attempt to do later on in this essay. The implicated notions of "experience," and of "explanatory dependence," need to be clarified (although it would be surprising if we were able to make the distinction so sharp that there are no borderline cases: no philosophical distinction is that precise).

While a certain amount of attention has been devoted to this project of clarifying the Distinction, much more attention has been devoted to explaining how, taking the cogency of the Distinction for granted, we could acquire the kinds of justification it distinguishes between. And while there are still many aspects of empirical knowledge that remain unclear, its existence is not in doubt; the real controversy centers on whether any belief is a priori justified.

In this connection, two notions have figured prominently: understanding and intuition. Some theories seek to explain all instances of a priori justification by showing that those propositions can be justifiably believed simply in virtue of being understood; others deploy the idea that we can know those propositions by intuiting their truth; others still, combine the two ideas (as, for example, in Boghossian 2020a, forthcoming). This project of explaining the source of a priori justification is far from complete or settled, although it would be fair to say that there has been much good work on the topic in recent times.[6]

Quine's well-known and highly influential skepticism about the a priori was fueled precisely by the claim that this explanatory project could not be satisfactorily carried out: there can be no good explanation for how we might be a priori justified in believing any particular proposition. As a result, Quine claimed, no belief could be a priori; all beliefs are, *au fond*, empirical.

In *Two Dogmas of Empiricism*, Quine picked up where the positivists had left off. He agreed with them that intuition was too obscure a notion to do any real explanatory work. That left understanding-based theories to contend with. The positivists had linked such understanding-based explanations of the a priori to the idea of *(metaphysical) analyticity* – some propositions could be known to be true in virtue of being understood *because* some propositions were "true purely in virtue of the concepts they contain and the way they are put together."[7] Unfortunately for the friends of the a priori, Quine argued, we can't make sense of this notion of analyticity. As a result, he concluded, given the prior rejection of intuitional theories, there was nothing left with which to explain how a priori justification was possible, and so no choice but to conclude that it is an empty notion.

Quine's influential argument suffers from several serious weaknesses. First, Quine leaves it utterly mysterious how we could have arrived at so much noncollusive convergence about our lists, if there is in fact nothing that distinguishes between them from an epistemological

[6] To mention just a few contributions: Bealer (2000), Bengson (2015), Chudnoff (2013), Malmgren (2006), and Warren (2020).

[7] Since Quine was writing in terms of sentences rather than propositions, his notion of metaphysical analyticity is the more familiar "true purely in virtue of meaning."

UNCORRECTED PROOFS

c08.indd 184 10-30-2023 17:43:17

point of view: surely, the fact that we are in such effortless agreement shows that we are onto *something*. Second, Quine's argument depends on finding fault with the notion of analyticity; but it is arguable that he had a good argument only against the *metaphysical* version of that notion, and not against the *epistemological* version, with the latter being the only one that is relevant to an explanation of the a priori (see Boghossian 1996, 2003, 2017). In addition, his unargued assumption that the notion of intuition can do no useful work is highly debatable (see some of the work cited above). Finally, the thoroughly empirical picture of all justification that Quine leaves us with can be shown to be hopeless, especially in connection with logic (see Boghossian and Wright, forthcoming), but also in connection with various other important domains (the normative, for example).

The implausibility of Quine's radical view helps explain the interest of the alternative form of skepticism about the distinction developed by Timothy Williamson (2007; Boghossian and Williamson 2020). Williamson very much wants to demote the distinction from the central place it has had in epistemology; but he wishes to do so while avoiding the Quinean embarrassment of having to say that all our agreement about how the distinction should be applied amounts to nothing. His basic idea is to concede that our noncollusive convergence shows that we are onto a real distinction while maintaining that that distinction is not an important one, not a "joint-carving" one for epistemology. In this respect, he says, it is analogous to the distinction between red and non-red fruits in botany: the distinction is obviously real, but of no great significance for the science of fruits (Williamson 2013: 291).

How does Williamson propose to show that the Distinction is real but unimportant? He has two interesting arguments for the claim.

The first attempts to argue that any theoretical attempt at formulating the Distinction precisely will end up being highly misaligned with the lists of examples as intuitively generated: paradigm examples of a priori belief will end up getting classified as a posteriori and vice versa. In a previous discussion (Boghossian 2020b), I argued that Williamson's case for this ambitious claim doesn't work: it rests on a spurious dilemma about how "experience" is to be understood. In this essay, I will try to go somewhat further and indicate (in Section B.2) how his challenge to formulate the Distinction in a way that aligns with our intuitions about cases might be met head-on.

Williamson's (2013) second and principal strategy for arguing for his claim, the one that has received the lion's share of attention, involves describing a pair of examples that are supposed to have the following features:

(A) One of them is a clear case of someone justifiably believing a proposition ("All crimson things are red") in an a priori way.
(B) The other is a clear case of someone justifiably believing another proposition ("All recent volumes of Who's Who are red") in an a posteriori way.
(C) The ways of justifiably believing involved in these examples are *broadly representative* of the ways in which propositions typically classified as a priori and a posteriori, respectively, are believed.
(D) There is no epistemologically interesting difference between the two ways of believing illustrated by these cases: they both rely on the same epistemic resources (acts of imagining an instance of a kind) used in the same way (to ground a generalization about the kind).

In a previous discussion (Boghossian 2020b), I argued that this argument fails because it is simply not true that you could know the bruited propositions in the ways described. Furthermore, those ways are not at all broadly representative, respectively, of a priori versus a posteriori ways of believing: indeed, so far as I know, no one has proposed *anything* like these

UNCORRECTED PROOFS

c08.indd 185 10-30-2023 17:43:17

accounts for knowledge of either proposition. There is, however, an additional problem with Williamson's argument that I haven't previously noted. Even if Williamson's premises were true, his argument would not establish the subtle conclusion that the distinction is real but unimportant, but only the cruder and much less palatable Quinean conclusion that the distinction is not real.

To see why, recall that the distinction between the a priori and the a posteriori is a distinction between two distinct *ways of justifiably believing* a proposition. That is to say, it is a theoretical distinction at the underlying epistemological level, not a distinction that can be verified observationally. Indeed, there is presumably no observable epistemological feature which all beliefs classified as "a priori" have and which all those classified as "a posteriori" lack. If our lists exemplify a distinction at all, they exemplify a distinction in underlying epistemological process.

Taken together, however, Williamson's claims (A) through (D) assert that there is no underlying epistemological difference between the central and broadly representative examples of those beliefs labeled "a priori" and those labeled "a posteriori." It is true, of course, that the *official* view is not that there is *no* difference, but that there is no "epistemologically interesting" difference between the two cases. However, one way in which two cases may fail to be interestingly different is for them not to be different at all. And that is in fact what is true of Williamson's cases, by his own admission:

> The problem is obvious. As characterized earlier, the cognitive processes underlying Norman's clearly a priori knowledge of (1) and his clearly a posteriori knowledge of (2) are almost exactly similar. If so, how can there be a deep epistemological difference between them? But if there is none, then the a priori–a posteriori distinction is epistemologically shallow.
>
> (Williamson 2013: 296–297)

The cases are not interestingly different because they are not different at all: they both involve an act of imagining an instance of a kind and generalizing from it. If Williamson's argument were sound, the only conclusion that would follow is that the distinction gestured at by our lists is illusory, not that it is real but unimportant.

In this respect, the distinction is not analogous to the red fruit/non-red fruit distinction in botany. In the case of the fruit, the stated distinction is observable, and there can be no question that it's real: we can tell by looking that there are red fruits and non-red fruits. We can then ask about this distinction whether it either is, or corresponds to, a distinction that is joint-carving for botany (and let's assume it's neither).

The a priori/a posteriori distinction, however, is not like that. There is no observable, ordinary-language distinction which the lists illustrate and about which we can ask whether it is, or corresponds to, an important distinction for epistemology. The distinction is stipulated to be a distinction at the underlying level. If there is no distinction at the underlying level, as Williamson claims through his examples, then there is no (exemplified) distinction at all. (There could still be a notional distinction, of course.)

Williamson likes to say that he's not denying the Distinction because he allows that there is a real distinction by "bottom-up" standards. There are two ways of reading this latter claim. On the first, natural reading, it is to say that our noncollusive convergence about which belief is to be labeled "a priori" and which is to be labeled "a posteriori" is strong *evidence* that the distinction is exemplified at the underlying epistemological level. On this reading, if Williamson's examples are accepted, the claim that there is a distinction by "bottom-up" standards would need to be ultimately rejected, even by him, since his examples argue that there is no such underlying distinction. On the alternative, much weaker, reading, to say that

UNCORRECTED PROOFS

c08.indd 186 10-30-2023 17:43:17

there is a distinction by "bottom-up" standards is just to reiterate the claim that there is noncollusive agreement on the labeling. On this reading, though, we don't *respect* the claim that the noncollusive agreement gives us reason to believe the distinction is real, since it's just a reiteration of the claim that there is such agreement – and that we already knew.

Williamson's central argument, then, cannot provide support for his distinctive brand of skepticism: it cannot show us how to demote the Distinction from its central place in epistemology while respecting the overwhelmingly plausible claim that our noncollusive convergence about the lists shows that we are onto something real.

B.2 Characterizing the Distinction

Let's turn to the question of how the distinction is to be characterized in such a way that it aligns sufficiently well with the intuitively generated lists. Obviously, this is a huge topic; I don't have the space here to consider everything that's been said, or the issues in all their complexity. My aim is merely to indicate the broad outlines of the proposal that I favor.

The way the notion of apriority is usually characterized is this:

> (Old) An a priori belief is a belief whose justification does not depend on experience.

For reasons that I will explain, I would prefer the following formulation:

> (New) An a priori belief is a belief whose justification is correctly explained without recourse to causal/tracking sources of information.[8]

A few remarks about the difference between the two formulations and why we ought to prefer New to Old.

"Correctly explained" instead of "depends on." As we have learned from a number of contexts, "depends on" is too crude for most purposes. In particular, the justification for a belief may *depend* on all kinds of factors that don't have anything to do with their explaining how it got its positive epistemic status. For example, since a belief couldn't have a justification without being a belief in the first place, whatever justification it has will *depend* on whatever experiences were required to have the belief. However, those enabling experiences won't explain the justification of the belief, but merely the belief itself. So, the explanatory criterion gives us a more nuanced tool with which to zero in on the relevant phenomenon.

"Causal/tracking sources of information" instead of "experience." Old and New agree in classifying as a posteriori all those beliefs whose justification depends on sense perception. But Old has trouble with *introspection*, which should intuitively also count as a source of a posteriori justification, and yet may not in every case involve experience (as, for example, when I know what my beliefs, understood as dispositional states, are). On the assumption that introspection, whether or not it involves experience, always involves somehow causal tracking, New gets this right.[9]

[8] This attempts to improve on the informal characterization given in Boghossian (2020b). It's very much in the spirit of Bonjour's (1998) proposal.

[9] There are some philosophers who have put forward noncausal accounts of introspective knowledge. For example, Burge (1988) has explored the idea that knowledge of one's own thoughts is achieved by thinking self-verifying thoughts of the form "With this very thought I am thinking that water is wet." For reasons explained in Boghossian (1989), I doubt that any such "logical" account of introspection could be correct. However, if it were true, it would no longer seem so implausible that self-knowledge is a priori. For a different approach see Wright (1989, 1998).

UNCORRECTED PROOFS

c08.indd 187 10-30-2023 17:43:17

Furthermore, *intuition* and *calculation* are both experiences, yet can count as sources of a priori justification. On the safe assumption that intuition and calculation don't work by causally tracking their subject matters, New gets this right while Old doesn't. Similarly for uses of the *imagination* which don't work by being constrained by our empirically supported background knowledge (see Boghossian 2020a).

Finally, if it were to turn out that we had some nonexperiential way of tracking distant events (*blindsight*, for example), that would not count intuitively as a source of a priori justification. Once again, Old would get this wrong, while New gets it right.

What about *memory*? Memory of past events is clearly a causal/tracking mechanism and so would count as a source of a posteriori knowledge. This gives rise to a well-known puzzle for friends of the a priori: since memory is going to be involved in any reasoning whatsoever, and since memory is a source of a posteriori knowledge, don't we have to admit that no belief arrived at on the basis of reasoning could be a priori? (see, e.g., Chisholm 1977). If we did concede that, though, that would wreak havoc with our lists, since many of the beliefs that appear on the a priori side could only have been arrived at inferentially.

Consider the following simple inference which I can run through completely "in my head" without the benefit of pencil and paper:

(Head)
If 17 is a prime number, then there is at least one prime that's smaller than 20.
17 is a prime number.
Therefore,
There is at least one prime that is smaller than 20.

It's very natural to think that not only can the premises of this inference be known a priori, so can the conclusion. But how is that possible given that I would have had to *remember* the premises in order to draw the conclusion?

In response, friends of the a priori appeal to a distinction between two kinds of memory – in Tyler Burge's (1993) terminology, between *substantive* memory and *preservative* memory. Substantive memory of a past event – as when I remember having dinner yesterday with friends – provides an a posteriori justification for belief in that past event. But the kind of memory that's operative in a short burst of reasoning like Head is not substantive in that sense. It does not involve recalling an event (thinking the premises, for example) and it does not provide any new justification for believing something not previously attained. Rather, it merely involves *preserving* the premises for use in reasoning. It is the buffer in which the premises linger before consciousness long enough to be operated on and reasoned with.[10]

Now, there are obvious psychological limits to how much can be in the buffer at any given time. No doubt this will vary somewhat from person to person, and even from time to time for a given person. But no matter how prodigious your short-term memory may be, there will be some inferences, or proofs, that are so long that you cannot carry them out solely in your head, but must write them down. (Suppose Euclid's proof that there is no greatest prime number is one of them.) Having written them down, you will need to look at them if you are to fully recall how the proof goes, or to check it for errors. So, even if we were to accept a distinction between substantive and preservative memory, and to insist that only preservative memory need be at work in reasoning, we would still face an issue about how the conclusion of a long proof could be a priori, given the seemingly ineliminable role for *perception* in arriving at the conclusion.

[10] As Burge notes, this distinction is reminiscent of one long recognized by psychologists between episodic memory and "semantic memory."

UNCORRECTED PROOFS

c08.indd 188 10-30-2023 17:43:17

Should we conclude that the limits of short-term memory set the limits to what can be known a priori? That would be implausible. Why is that? The reason isn't simply that we have already committed ourselves to Euclid's proof being typically known a priori in our intuitive classifications ("the bottom-up standards"). It's because, even if you reversed yourself on that classification, you would still want to mark the fact, obviously hugely significant for epistemology, that Euclid's proof *would be known a priori even by you*, if you simply had a somewhat enhanced capacity for short-term memory, while no amount of idealizing that capacity would render "It's sunny now" a priori knowable by you. Clearly, the capacity of short-term memory is not a significant parameter as far as the distinction is concerned.

And we can accommodate that insight by emphasizing that perception plays two very different roles in the two cases. In the case of the latter, perception is needed to explain how you are justified in believing that it is sunny outside now; in the case of the latter, perception is needed only in order to explain how you have *access* to the justified premises and inferential steps of a long proof in terms of which the justification for the conclusion is to be explained.

And so, we see that, just as we always knew, a proper understanding of the a priori/a posteriori distinction depends on a distinction between a justification-*enabling* versus a justification-*explaining* use of sense experience, although as we can also see, that distinction has to be understood in a more capacious sense than has been traditional. You may need sense experience to so much as be able to think the proposition "If it's sunny then it's sunny" (you might need it to acquire the concept "sunny.") You may also need sense experience to access that proposition (along with its judgmental force and positive epistemic status) in a long chain of reasoning that you cannot keep in your head, but which you need to write down. However, once all of those access conditions have been met, some propositions will be such that your justification for believing them will be explicable solely by adverting to sense experience, and others won't.[11]

The central project for friends of the a priori is to develop worked-out theories of where nonempirical justification comes from and how it works. What role does the understanding play? And what role, if any, does intuition have? And how is it possible to gain knowledge of the world by such means? These are difficult and still contested questions; but, as I emphasized earlier, there is an increasingly sophisticated body of literature on them and much progress has been made. Here is not the place to further delve into these tricky matters. I will devote the remainder of this essay to exploring the question why it should matter whether there is a distinction between the a priori and the a posteriori along the lines that I've described.

B.3 Why Does the Distinction Matter?

Why should it matter whether we accept that there are beliefs whose justification can be explained without recourse to causal/tracking sources of information? It matters not only because the claim seems true, but because it is impossible to see how to give a satisfactory account of the justification that we have for some of our most important beliefs, or rules of inference, in terms of such sources of information.

Quine understood very well the extraordinarily radical consequences of his rejection of the distinction. He understood that it implied that everything, including the most basic logical truth or rule of inference, could only count as epistemically justified if it earned its place in an

[11] For an argument that we may need a yet more capacious version of the distinction between enabling and justifying, see Russell (2022).

UNCORRECTED PROOFS

c08.indd 189 10-30-2023 17:43:17

overall empirical theory of the world that was doing better (by given metrics) than the available alternatives. And he also understood, and embraced, the consequence that, in principle, even the most basic logical rule might get overturned, if rejecting or restricting it proved to be the best route to developing overall scientific theory. As Quine famously put it:

> Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system. Even a statement very close to the periphery can be held true in the face of recalcitrant experience by pleading hallucination or by amending certain statements of the kind called logical laws. Conversely, by the same token, no statement is immune to revision. Revision even of the logical law of the excluded middle has been proposed as a means of simplifying quantum mechanics; and what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle?
>
> (Quine 1953: 43)

This is a shocking view, although we have becomes so used to it that its shock value has faded. The idea that such basic logical truths as

> If *p*, then *p*
> If *p* and (if *p*, then *q*), then *q*
> If all Fs are Gs, then this F is a G

might be shown to be false because rejecting them might "simplify the formulation of quantum mechanics" is a jaw-dropping view. This is so for many reasons, some of which leap out at you, others of which may require more explanation. Among those reasons we may cite, first: because it would be very hard to understand what someone could be so much as *saying* if they denied that if *p*, then *p*. Second, because it's unclear how any such view could accommodate a feature of logic that would appear to be nonoptional, namely, that logical laws apply in *all* worlds, and not just in nomologically accessible ones. If the truth of a logical law is so tied up with which theories best describe how the *actual* world is, how could we be so confident that it would apply even in counter-nomological worlds?[12] And, finally (for now), because it's hard to see how *any* theoretical reasoning could take place – including the kind of reasoning that would be required in thinking about how it would be best to revise logic on a Quinean picture – that didn't involve *some version* of universal instantiation, modus ponens or the law of noncontradiction; if that's right, though, these basic logical rules are not up for rejection in the way that the Quinean picture implies (for more details see Boghossian and Wright, forthcoming).

Skeptics about the a priori like to complain that its friends have yet to come up with a completely satisfying account of how a priori justification is attained. They forget that they have not even come close to explaining how we are to avoid the absurdities incurred by the view that everything, including basic logic, is to be treated on a par with empirical belief.

Their problem is not confined to the case of logic. Empirical knowledge of fundamental normative principles would present just as much of a challenge. Take, for example, a principle that plausibly governs the justification we receive from visual perception:

> (Perception) If it perceptually seems to you that *p*, you are prima facie justified in believing that *p*.[13]

[12] A point I owe to Crispin Wright. The point generalizes to any necessary truth.

[13] If you don't think that is the correct formulation of the principle governing perceptual justification, substitute your own preferred version; it won't matter to the point being made.

UNCORRECTED PROOFS

c08.indd 190 10-30-2023 17:43:17

What possible experimental evidence could strengthen or decrease your credence in this principle? Any such evidence would itself have to be perceptual; so, you would have to rely on it (or an equivalent principle) in order to make use of that evidence. To then use that evidence to strengthen your credence in Perception would involve an illegitimate type of bootstrapping. To use that evidence to defeat the principle would appear to involve a type of incoherence. The examples could be multiplied.

There are many beliefs – indeed, some of the most important beliefs we have – whose justification is attained, just as it has always seemed, through reason alone, without the benefit of our ability to causally track the contingent properties of our (inner or outer) environment. It has been easy to mount skeptical attacks on this claim because it has turned out to be exceptionally difficult to give a fully satisfying account of how that justification is achieved. However, the claim has been able to withstand these attacks not only because it seems overwhelmingly true, but because denying it in full generality courts incoherence. It is time to face up to the fact that the a priori is here to stay.[14]

C. Response to Boghossian

Timothy Williamson

1. Paul Boghossian and I have been debating a priori knowledge and justification since 2003, as recorded in our book *Debating the A Priori* (Boghossian and Williamson 2020). Some of his opening essay here summarizes parts of those exchanges, though too briefly to do justice to either side's arguments. For reasons of space, I will not go over that old ground here; the interested reader can consult that book to see how I answer his objections. This response concentrates on the newer parts of his essay.

The terms "a priori" and "a posteriori" occur frequently in philosophical discourse, but without agreement on how they should be defined. They are applicable to specific cases of knowledge, which is fundamental for me, and justification, which is fundamental for Boghossian. So used, they have nonempty, nonoverlapping extensions. Although various complications arise, they do not matter for present purposes. So there is a real distinction between the a priori and the a posteriori. However, I argue, the distinction, though nonillusory, is *superficial*; it does not cut at an epistemological joint.

Boghossian claims that I have mischaracterized the upshot of my own arguments because "The distinction is stipulated to be a distinction at the underlying level." Thus, if there is no underlying distinction, there is no distinction at all. His reason for that claim is that if we draw up lists of agreed examples of the a priori and of the a posteriori, "There is no observable, ordinary-language distinction which the lists illustrate and about which we can ask whether it is, or corresponds to, an important distinction for epistemology." But that argument underestimates the variety of ways in which meanings are determined. Consider the vocabulary we have inherited from the ancient theory of the four humors for classifying types of human character: "sanguine" (blood), "choleric" (yellow bile), "melancholy" (black bile), "phlegmatic" (phlegm). They are not observation terms. They are used in educated English, but so are "a priori" and "a posteriori." Speakers familiar with them may well be in significant noncollusive agreement about their application to particular cases. The terms are

[14] For helpful comments on an earlier version of this essay, I am grateful to Yu Guo, Declan Smithies, Gideon Rosen, and Crispin Wright.

UNCORRECTED PROOFS

c08.indd 191 10-30-2023 17:43:17