

on more than we like to admit). But the fact that the canons of rationality are themselves evolving doesn't mean they don't exist (*pace* Feyerabend, *pace* Foucault!), nor does it mean that, in the course of the evolution, *anything whatsoever* (including acceptance of the absolutely inconsistent rule) might occur. Evolution, in the domain of *instruments*, does not imply total, protean, lack of definite structure.

But, after all, just how important is it that Quine is wrong in his *total* rejection of the *a priori*? In one way it is not very important. We do not have a good *theory* of rationality, and are unlikely to have one in the foreseeable future. Lacking the 'rigid designator' of rationality, the theoretical definition which tells us what rationality is in every possible world (as 'water is H<sub>2</sub>O' tells us what water is in every possible world), it is virtually

hopeless to show with any semblance of good argument that any specific statement is such that it would be irrational to ever give it up (apart from special examples, such as the one I constructed). Nor do we really need a proof that a statement is *a priori* in this sense (rationally unrevisable) very often. If a statement has the property that *we cannot now describe* any circumstances under which it would be rational to give it up, that will surely suffice for most purposes of philosophical argument. But, if it is always dangerous to take on the burden of trying to show that a statement is absolutely *a priori*, the foregoing reflections show that it is not just dangerous but actually wrong to make the quick leap from the fact that it is dangerous to claim that any statement is absolutely *a priori* to the absolute claim that there are no *a priori* truths.

## A Priori Knowledge

PHILIP KITCHER

### I

"A priori" has been a popular term with philosophers at least since Kant distinguished between a priori and a posteriori knowledge. Yet, despite the frequency with which it has been used in twentieth century philosophy, there has been little discussion of the concept of apriority.<sup>1</sup> Some writers seem to take it for granted that there are propositions, such as the truths of logic and mathematics, which are a priori; others deny that there

are any a priori propositions. In the absence of a clear characterization of the a priori/a posteriori distinction, it is by no means obvious what is being asserted or what is being denied.

"A priori" is an epistemological predicate. What is *primarily* a priori is an item of knowledge.<sup>2</sup> Of course, we can introduce a derivative use of "a priori" as a predicate of propositions:<sup>3</sup> a priori propositions are those which we could know a priori. Somebody might protest that current practice is to define the notion of an a priori proposition outright, by taking the class of a priori propositions to consist of the truths of logic and

Reprinted from *The Philosophical Review*, LXXVI (1980), 3–23, with permission of Philip Kitcher and publisher, copyright 1980 by *The Philosophical Review*.

mathematics (for example). But when philosophers allege that truths of logic and mathematics are a priori, they do not intend merely to recapitulate the definition of a priori propositions. Their aim is to advance a thesis about the epistemological status of logic and mathematics.

To understand the nature of such epistemological claims, we should return to Kant, who provided the most explicit characterization of a priori knowledge: "we shall understand by a priori knowledge, not knowledge which is independent of this or that experience, but knowledge absolutely independent of all experience."<sup>4</sup> While acknowledging that Kant's formulation sums up the classical notion of apriority, several recent writers who have discussed the topic have despaired of making sense of it.<sup>5</sup> I shall try to show that Kant's definition can be clarified, and that the concept of a priori knowledge can be embedded in a naturalistic epistemology.

## II

Two questions naturally arise. What are we to understand by "experience"? And what is to be made of the idea of independence from experience? Apparently, there are easy answers. Count as a person's experience the stream of her sensory encounters with the world, where this includes both "outer experience," that is, sensory states caused by stimuli external to the body, and "inner experience," that is, those sensory states brought about by internal stimuli. Now we might propose that someone's knowledge is independent of her experience just in case she could have had that knowledge whatever experience she had had. To this obvious suggestion there is an equally obvious objection. The apriorist is not ipso facto a believer in innate knowledge: indeed, Kant emphasized the difference between the two types of knowledge. So we cannot accept an analysis which implies that a priori knowledge could have been obtained given minimal experiences.<sup>6</sup>

Many philosophers (Kant included) contend both that analytic truths can be known a priori and that some analytic truths involve concepts which could only be acquired if we were to have particular kinds of experience. If we are to defend their doctrines from immediate rejection, we must allow a minimal role to experience, even in a priori knowledge. Experience may be needed to provide some concepts. So we might modify our proposal: knowledge is independent of experience if any experience which would enable us to acquire the concepts involved would enable us to have the knowledge.

It is worth noting explicitly that we are concerned here with the *total* experience of the knower. Suppose that you acquire some knowledge empirically. Later you deduce some consequences of this empirical knowledge. We should reject the suggestion that your knowledge of those consequences is independent of experience because, at the time you perform the deduction, you are engaging in a process of reasoning which is independent of the sensations you are then having.<sup>7</sup> As Kant recognized,<sup>8</sup> your knowledge, in cases like this, is dependent on your total experience: different total sequences of sensations would not have given you the premises for your deductions.

Let us put together the points which have been made so far. A person's experience at a particular time will be identified with his sensory state at the time. (Such states are best regarded physicalistically in terms of stimulation of sensory receptors, but we should recognize that there are both "outer" and "inner" receptors.) The total sequence of experiences X has had up to time *t* is *X's life at t*. A life will be said to be *sufficient for X for p* just in case X could have had that life and gained sufficient understanding to believe that *p*. (I postpone, for the moment, questions about the nature of the modality involved here.) Our discussion above suggests the use of these notions in the analysis of a priori knowledge: X knows a priori that *p* if and only if X knows that *p* and, given any life sufficient for X for *p*, X could have had

that life and still have known that *p*. Making temporal references explicit: at time *t* *X* knows a priori that *p* just in case, at time *t*, *X* knows that *p* and, given any life sufficient for *X* for *p*, *X* could have had that life at *t* and still have known, at *t*, that *p*. In subsequent discussions I shall usually leave the temporal references implicit.

Unfortunately, the proposed analysis will not do. A clear-headed apriorist should admit that people can have empirical knowledge of propositions which can be known a priori. However, on the account I have given, if somebody knows that *p* and if it is possible for her to know a priori that *p*, then, apparently, given any sufficiently rich life she could know that *p*, so that she would meet the conditions for a priori knowledge that *p*. (This presupposes that modalities "collapse," but I don't think the problem can be solved simply by denying the presupposition.) Hence it seems that my account will not allow for empirical knowledge of propositions that can be known a priori.

We need to amend the analysis. We must differentiate situations in which a person knows something empirically which could have been known a priori from situations of actual a priori knowledge. The remedy is obvious. What sets apart corresponding situations of the two types is a difference in the ways in which what is known is known. An analysis of a priori knowledge must probe the notion of knowledge more deeply than we have done so far.

### III

We do not need a general analysis of knowledge, but we do need the *form* of such an analysis. I shall adopt an approach which may appropriately be called "the psychological account of knowledge."<sup>9</sup> The root idea is that the question of whether a person's true belief counts as knowledge depends on whether the presence of that true belief can be explained in an appropriate fashion. The difference between an item of knowledge

and mere true belief turns on the factors which produced the belief; thus the issue revolves around the way in which a particular mental state was generated. It is important to emphasize that, at different times, a person may have states of belief with the same content, and these states may be produced by different processes. The claim that a process produces a belief is to be understood as the assertion that the presence of the current state of belief is to be explained through a description of that process. Hence the account is not committed to supposing that the original formation of a belief is relevant to the epistemological status of later states of belief in the same proposition.<sup>10</sup>

The question of what conditions must be met if a belief is to be explained in an appropriate fashion is central to epistemology, but it need not concern us here. My thesis is that the distinction between knowledge and true belief depends on the characteristics of the process which generates the belief, and this thesis is independent of specific proposals about what characteristics are crucial. Introducing a useful term, let us say that some processes *warrant* the beliefs they produce, and that these processes are *warrants* for such beliefs. The general view of knowledge I have adopted can be recast as the thesis that *X* knows that *p* just in case *X* correctly believes that *p* and *X*'s belief was produced by a process which is a warrant for it. Leaving the task of specifying the conditions on warrants to general epistemology, my aim is to distinguish a priori knowledge from a posteriori knowledge. We discovered above that the distinction requires us to consider the ways in which what is known is known. Hence I propose to reformulate the problem: let us say that *X* knows a priori that *p* just in case *X* has a true belief that *p* and that belief was produced by a process which is an *a priori warrant* for it. Now the crucial notion is that of an *a priori warrant*, and our task becomes that of specifying the conditions which distinguish *a priori warrants* from other warrants.

At this stage, some examples may help us to see how to draw the distinction. Perception is an obvious type of process which philosophers have supposed *not* to engender a priori knowledge. Putative a priori warrants are more controversial. I shall use Kant's notion of pure intuition as an example. This is not to endorse the claim that processes of pure intuition are a priori warrants, but only to see what features of such processes have prompted Kant (and others) to differentiate them from perceptual processes.

On Kant's theory, processes of pure intuition are supposed to yield a priori mathematical knowledge. Let us focus on a simple geometrical example. We are supposed to gain a priori knowledge of the elementary properties of triangles by using our grasp on the concept of triangle to construct a mental picture of a triangle and by inspecting this picture with the mind's eye.<sup>11</sup> What are the characteristics of this kind of process which make Kant want to say that it produces knowledge which is independent of experience? I believe that Kant's account implies that three conditions should be met. The same type of process must be *available* independently of experience. It must produce *warranted* belief independently of experience. And it must produce *true* belief independently of experience. Let us consider these conditions in turn.

According to the Kantian story, if our life were to enable us to acquire the appropriate concepts (the concept of a triangle and the other geometrical concepts involved) then the appropriate kind of pure intuition would be available to us. We could represent a triangle to ourselves, inspect it, and so reach the same beliefs. But, if the process is to generate *knowledge* independently of experience, Kant must require more of it. Given any sufficiently rich life, if we were to undergo the same type of process and gain the same beliefs, then those beliefs would be warranted by the process. Let us dramatize the point by imagining that experience is unkind. Suppose that we are presented with experiments which are cunningly contrived so as to make it appear that some of our basic

geometrical beliefs are false. Kant's theory of geometrical knowledge presupposes that if, in the circumstances envisaged, a process of pure intuition were to produce geometrical belief then it would produce warranted belief, despite the background of misleading experience.

So far I have considered how a Kantian process of pure intuition might produce warranted belief independently of experience. But to generate *knowledge* independently of experience, a priori warrants must produce warranted *true* belief in counterfactual situations where experiences are different. This point does not emerge clearly in the Kantian case because the propositions which are alleged to be known a priori are taken to be necessary, so that the question of whether it would be possible to have an a priori warrant for a false belief does not arise. Plainly, we could ensure that a priori warrants produce warranted *true* belief independently of experience by declaring that a priori warrants only warrant necessary truths. But this proposal is unnecessarily strong. Our goal is to construe a priori knowledge as knowledge which is independent of experience, and this can be achieved, without closing the case against the contingent a priori, by supposing that, in a counterfactual situation in which an a priori warrant produces belief that *p* then *p*. On this account, a priori warrants are ultra-reliable; they never lead us astray.<sup>12</sup>

Summarizing the conditions that have been uncovered, I propose the following analysis of a priori knowledge.

- (1) X knows a priori that *p* if and only if X knows that *p* and X's belief that *p* was produced by a process which is an a priori warrant for it.
- (2)  $\alpha$  is an a priori warrant for X's belief that *p* if and only if  $\alpha$  is a process such that, given any life *e*, sufficient for X for *p*, then
  - (a) some process of the same type could produce in X a belief that *p*
  - (b) if a process of the same type were to produce in X a belief that *p* then it would warrant X in believing that *p*
  - (c) if a process of the same type were to produce in X a belief that *p* then *p*.

It should be clear that this analysis yields the desired result that, if a person knows a priori that  $p$  then she could know that  $p$  whatever (sufficiently rich) experience she had had. But it goes beyond the proposal of §II in spelling out the idea that the knowledge be obtainable in the same way. Hence we can distinguish cases of empirical knowledge of propositions which could be known a priori from cases of actual a priori knowledge.

#### IV

In this section, I want to be more explicit about the notion of "types of processes" which I have employed, and about the modal and conditional notions which figure in my analysis. To specify a process which produces a belief is to pick out some terminal segment of the causal ancestry of the belief. I think that, without loss of generality, we can restrict our attention to those segments which consist solely of states and events internal to the believer.<sup>13</sup> Tracing the causal ancestry of a belief beyond the believer would identify processes which would not be available independently of experience, so that they would violate our conditions on a priori warrants.

Given that we need only consider psychological processes, the next question which arises is how we divide processes into types. It may seem that the problem can be side-stepped: can't we simply propose that to defend the apriority of an item of knowledge is to claim that that knowledge was produced by a psychological process and that *that very process* would be available and would produce warranted true belief in counterfactual situations where experience is different? I think it is easy to see how to use this proposal to rewrite (2) in a way which avoids reference to "types of processes." I have not adopted this approach because I think that it short-cuts important questions about what makes a process the same in different counterfactual situations.

Our talk of processes which produce belief was originally introduced to articulate

the idea that some items of knowledge are obtained in the same way while others are obtained in different ways. To return to our example, knowing a theorem on the basis of hearing a lecture and knowing the same theorem by following a proof count, intuitively, as different ways of knowing the theorem. Our intuitions about this example, and others, involve a number of different principles of classification, with different principles appearing in different cases. We seem to divide belief-forming processes into types by considering content of beliefs, inferential connections, causal connections, use of perceptual mechanisms and so forth. I suggest that these principles of classification probably do not give rise to one definite taxonomy, but that, by using them singly, or in combination, we obtain a number of different taxonomies which we can and do employ. Moreover, within each taxonomy, we can specify types of processes more or less narrowly.<sup>14</sup> Faced with such variety, what characterization should we pick?

There is probably no privileged way of dividing processes into types. This is not to say that our standard principles of classification will allow *anything* to count as a type. Somebody who proposed that the process of listening to a lecture (or the terminal segment of it which consists of psychological states and events) belongs to a type which consists of itself and instances of following a proof, would flout *all* our principles for dividing processes into types. Hence, while we may have many admissible notions of types of belief-forming processes, corresponding to different principles of classification, some collections of processes contravene all such principles, and these cannot be admitted as genuine types.<sup>15</sup>

My analysis can be read as issuing a challenge to the apriorist. If someone wishes to claim that a particular belief is an item of a priori knowledge then he must specify a segment of the causal ancestry of the belief, consisting of states and events internal to the believer, and type-identity conditions which conform to some principle (or set of principles) of classification which are standardly

employed in our divisions of belief-forming processes (of which the principles I have indicated above furnish the most obvious examples). If he succeeds in doing this so that the requirements in (2) are met, his claim is sustained; if he cannot, then his claim is defeated.

The final issue which requires discussion in this section is that of explaining the modal and conditional notions I have used. There are all kinds of possibility, and claims about what is possible bear an implicit relativization to a set of facts which are held constant.<sup>16</sup> When we say, in (2), that, given any sufficiently rich life, X could have had a belief which was the product of a particular type of process, should we conceive of this as merely logical possibility or are there some features of the actual world which are tacitly regarded as fixed? I suggest that we are not just envisaging any logically possible world. We imagine a world in which X has similar mental powers to those he has in the actual world. By hypothesis, X's experience is different. Yet the capacities for thinking, reasoning, and acquiring knowledge which X possesses as a member of *homo sapiens* are to remain unaffected: we want to say that X, *with the kinds of cognitive capacities distinctive of humans*, could have undergone processes of the appropriate type, even if his experiences had been different.<sup>17</sup>

Humans might have had more faculties for acquiring knowledge than they actually have. For example, we might have had some strange ability to "see" what happens on the other side of the Earth. When we consider the status of a particular type of process as an a priori warrant, the existence of worlds in which such extra faculties come into play is entirely irrelevant. Our investigation focuses on the question of whether a particular type of process would be available to a person with the kinds of faculties people actually have, not on whether such processes would be available to creatures whose capacities for acquiring knowledge are augmented or diminished. Conditions (2(b)) and (2(c)) are to be read in similar fashion. Rewriting (2(b)) to make the form of the conditional

explicit, we obtain: for any life e sufficient for X for p and for any world in which X has e, in which he believes that p, in which his belief is the product of a process of the appropriate kind, and *in which X has the cognitive capacities distinctive of humans*, X is warranted in believing that p. Similarly, (2(c)) becomes: for any life e sufficient for X for p and for any world in which X has e, in which he believes that p, in which his belief is the product of a process of the appropriate kind, and *in which X has the cognitive capacities distinctive of humans*, p. Finally, the notion of a life's being sufficient for X for p also bears an implicit reference to X's native powers. To say that a particular life enables X to form certain concepts is to maintain that, given the genetic programming with which X is endowed, that life allows for the formation of the concepts.

The account I have offered can be presented more graphically in the following way. Consider a human as a cognitive device, endowed initially with a particular kind of structure. Sensory experience is fed into the device and, as a result, the device forms certain concepts. For any proposition p, the class of experiences which are sufficiently rich for p consists of those experiences which enable the device, with the kind of structure it actually has, to acquire the concepts to believe that p. To decide whether or not a particular item of knowledge that p is an item of a priori knowledge we consider whether the type of process which produced the belief that p is a process which would have been available to the device, with the kind of structure it actually has, if different sufficiently rich experiences had been fed into it, whether, under such circumstances, processes of the type would warrant belief that p, and would produce true belief that p.

Seen in this way, claims about apriority are implicitly indexical, in that they inherit the indexical features of "actual."<sup>18</sup> If this is not recognized, use of "a priori" in modal contexts can engender confusion. The truth value of "Possibly, X knows a priori that p" can be determined in one of two ways: we

may consider the proposition expressed by the sentence at our world, and inquire whether there is a world at which that proposition is true; or we may ask whether there is a world at which the sentence expresses a true proposition. Because of the covert indexicality of "a priori," these lines of investigation may yield different answers. I suspect that failure to appreciate this point has caused trouble in assessing theses about the limits of the a priori. However, I shall not pursue the point here.<sup>19</sup>

## V

At this point, I want to address worries that my analysis is too liberal, because it allows some of our knowledge of ourselves and our states to count as a priori. Given its Kantian psychological underpinnings, the theory appears to favor claims that some of our self-knowledge is a priori. However, two points should be kept in mind. Firstly, the analysis I have proposed can only be applied to cases in which we know enough about the ways in which our beliefs are warranted to decide whether or not the conditions of (2) are met. In some cases, our lack of a detailed account of how our beliefs are generated may mean that no firm decision about the apriority of an item of knowledge can be reached. Secondly, there may be cases, including cases of self-knowledge, in which we have no clear pre-analytic intuitions about whether a piece of knowledge is a priori.

Nevertheless, there are some clear cases. Obviously, any theory which implied that I can know a priori that I am seeing red (when, in fact, I am) would be suspect. But, when we apply my analysis, the unwanted conclusion does not follow. For, if the process which leads me to believe that I am seeing red (when I am) can be triggered in the absence of red, then (2(c)) would be violated. If the process cannot be triggered in the absence of red, then, given some sufficiently rich experiences, the process will not be available, so that (2(a)) will be violated. In general, knowledge of any involuntary

mental state—such as pains, itches or hallucinations—will work in the same way. Either the process which leads from the occurrence of pain to the belief that I am in pain can be triggered in the absence of pain, or not: if it can, (2(c)) would be violated, if it cannot, then (2(a)) would be violated.

This line of argument can be sidestepped when we turn to cases in which we have the power, independently of experience, to put ourselves into the appropriate states. For, in such cases, one can propose that the processes which give us knowledge of the state cannot be triggered in the absence of the states themselves *and* that the processes are always available because we can always put ourselves into the states.<sup>20</sup> On this basis, we might try to conclude that we have a priori knowledge that we are imagining red (when we are) or thinking of Ann Arbor (when we are). However, the fact that such cases do not fall victim to the argument of the last paragraph does not mean that we are compelled to view them as cases of a priori knowledge. In the first place, the thesis that the processes through which we come to know our imaginative feats and our voluntary thoughts cannot be triggered in the absence of the states themselves requires evaluation—and, lacking detailed knowledge of those processes, we cannot arrive at a firm judgment here. Secondly, the processes in question will be required to meet (2(b)) if they are to be certified as a priori warrants. This means that, whatever experience hurls at us, beliefs produced by such processes will be warranted. We can cast doubt on this idea by imagining that our experience consists of a lengthy, and apparently reliable, training in neurophysiology, concluding with a presentation to ourselves of our own neurophysiological organization which appears to show that our detection of our imaginative states (say) is slightly defective, that we always make mistakes about the contents of our imaginings. If this type of story can be developed, then (2(b)) will be violated, and the knowledge in question will not count as a priori. But, even if it cannot be coherently extended, and even if my analysis



does judge our knowledge of states of imagination (and other "voluntary" states) to be a priori, it is not clear to me that this consequence is counterintuitive.

In fact, I think that one can make a powerful case for supposing that *some* self-knowledge is a priori. At most, if not all, of our waking moments, each of us knows of herself that she exists.<sup>21</sup> Although traditional ideas to the effect that self-knowledge is produced by some "non-optical inner look" are clearly inadequate, I think it is plausible to maintain that there are processes which do warrant us in believing that we exist—processes of reflective thought, for example—and which belong to a general type whose members would be available to us independently of experience.<sup>22</sup> Trivially, when any such process produces in a person a belief that she exists that belief is true. All that remains, therefore, is to ask if the processes of the type in question inevitably warrant belief in our own existence, or whether they would fail to do so, given a suitably exotic background experience. It is difficult to settle this issue conclusively without a thorough survey of the ways in which reflective belief in one's existence can be challenged by experience, but perhaps there are Cartesian grounds for holding that, so long as the belief is the product of reflective thought, the believer is warranted, no matter how wild his experience may have been. If this is correct, then at least some of our self-knowledge will be a priori. However, in cases like this, attributions of apriority seem even less vulnerable to the criticism that they are obviously incorrect.

At this point we must consider a doctrinaire objection. If the conclusion of the last paragraph is upheld then we can know some contingent propositions a priori.<sup>23</sup> Frequently, however, it is maintained that only necessary truths can be known a priori. Behind this contention stands a popular argument.<sup>24</sup> Assume that a person knows a priori that *p*. His knowledge is independent of his experience. Hence he can know that *p* without any information about the kind of world he inhabits. So, necessarily *p*.

This hazy line of reasoning rests on an intuition which is captured in the analysis given above. The intuition is that a priori warrants must be ultra-reliable: if a person is entitled to ignore empirical information about the type of world she inhabits then that must be because she has at her disposal a method of arriving at belief which guarantees *true* belief. (This intuition can be defended by pointing out that if a method which could produce false belief were allowed to override experience, then we might be blocked from obtaining knowledge which we might otherwise have gained.) In my analysis, the intuition appears as (2(c)).<sup>25</sup>

However, when we try to clarify the popular argument we see that it contains an invalid step. Presenting it as a *reductio*, we obtain the following line of reasoning. Assume that a person knows a priori that *p* but that it is not necessary that *p*. Because *p* is contingent there are worlds at which *p* is false. Suppose that the person had inhabited such a world and behaved as she does at the actual world. Then she would have had an a priori warrant for a false belief. This is debarred by (2(c)). So we must conclude that the initial supposition is erroneous: if someone really does know a priori that *p* then *p* is necessary.

Spelled out in this way, the argument fails. We are not entitled to conclude from the premise that there are worlds at which *p* is false the thesis that there are worlds at which *p* is false *and* at which the person behaves as she does at the actual world. There are a number of propositions which, although they could be false, could not both be false and also believed by us. More generally, there are propositions which could not both be false and also believed by us in particular, definite ways. Obvious examples are propositions about ourselves and their logical consequences: such propositions as those expressed by tokens of the sentences "I exist," "I have some beliefs," "There are thoughts," and so forth. Hence the attempted *reductio* breaks down and allows for the possibility of a priori knowledge of some contingent propositions.



I conclude that my analysis is innocent of the charge of being too liberal in ascribing to us a priori knowledge of propositions about ourselves. Although it is plausible to hold that my account construes some of our self-knowledge as a priori, none of the self-knowledge it takes to be a priori is clearly empirical. Moreover, it shows how a popular argument against the contingent a priori is flawed, and how certain types of contingent propositions—most notably propositions about ourselves—escape that argument. Thus I suggest that the analysis illuminates an area of traditional dispute.

## VI

I now want to consider two different objections to my analysis. My replies to these objections will show how the approach I have developed can be further refined and extended.

The first objection, like those considered above, charges that the analysis is too liberal. My account apparently allows for the possibility that a priori knowledge could be gained through perception. We can imagine that some propositions are true at any world of which we can have experience, and that, given sufficient experience to entertain those propositions, we could always come to know them on the basis of perception. Promising examples are the proposition that there are objects, the proposition that some objects have shapes, and other, similar propositions. In these cases, one can argue that we cannot experience worlds at which they are false and that any (sufficiently rich) experience would provide perceptual warrant for belief in the propositions, regardless of the specific content of our perceptions. If these points are correct (and I shall concede them both, for the sake of argument), then perceptual processes would qualify as a priori warrants. Given any sufficiently rich experience, some perceptual process would be available to us, would produce warranted belief and, *ex hypothesi*, would produce warranted *true* belief.

Let us call cases of the type envisaged

cases of *universally empirical* knowledge. The objection to my account is that it incorrectly classifies universally empirical knowledge as a priori knowledge. My response is that the classical notion of apriority is too vague to decide such cases: rather, this type of knowledge only becomes apparent when the classical notion is articulated. One could defend the classification of universally empirical knowledge as a priori by pointing out that such knowledge requires no particular type of experience (beyond that needed to obtain the concepts, of course). One could oppose that classification by pointing out that, even though the content of the experience is immaterial, the knowledge is still gained by perceiving, so that it should count as a posteriori.

If the second response should seem attractive, it can easily be accommodated by recognizing a stronger and a weaker notion of apriority. The weaker notion is captured in (1) and (2). The stronger adds an extra requirement: no process which involves the operation of a perceptual mechanism is to count as an a priori warrant.

At this point, it is natural to protest that the new condition makes the prior analysis irrelevant. Why not define a priori knowledge outright as knowledge which is produced by processes which do not involve perceptual mechanisms? The answer is that the prior conditions are not redundant: knowledge which is produced by a process which does not involve perceptual mechanisms need not be independent of experience. For the process may fail to generate warranted belief against a backdrop of misleading experience. (Nor may it generate true belief in all relevant counterfactual situations.) So, for example, certain kinds of thought-experiments may generate items of knowledge given a particular type of experience, but may not be able to sustain that knowledge against misleading experiences. Hence, if we choose to exclude universally empirical knowledge from the realm of the a priori in the way suggested, we are building on the analysis given in (1) and (2), rather than replacing it.

A different kind of criticism of my analysis is to accuse it of revealing the emptiness of the classical notion of apriority. Someone may suggest that, in exposing the constraints on a priori knowledge, I have shown that there could be very little a priori knowledge. Although I believe that this suggestion is incorrect, it is worth pointing out that, even if it is granted, my approach allows for the development of weaker notions which may prove epistemologically useful.

Let me first note that we can introduce approximations to a priori knowledge. Suppose that A is any type of process all of whose instances culminate in belief that p. Define the *supporting class* of A to be that class of lives, e, such that, (a) given e, some process in A could occur (and so produce belief that p), (b) given e, any process in A which occurred would produce warranted true belief that p. (Intuitively, the supporting class consists of those lives which enable processes of the type in question to produce knowledge.) The *defeating class* of A is the complement of the supporting class of A within the class of lives which are sufficient for p. A priori warrants are those processes which belong to a type whose defeating class is null. But we can be more liberal, and allow approximations to a priori knowledge by considering the size and/or nature of the defeating class. We might, for example, permit the defeating class to contain those radically disruptive experiences beloved of sceptics. Or we can define a notion of *contextual* apriority by allowing the defeating class to include experiences which undermine "framework principles."<sup>26</sup> Or we may employ a concept of *comparative* apriority by ordering defeating classes according to inclusion relations. Each of these notions can serve a useful function in delineating the structure of our knowledge.

## VII

Finally, I want to address a systematic objection to my analysis. The approach I have taken is blatantly psychologistic. Some phi-

losophers may regard these psychological complications as objectionable intrusions into epistemology. So I shall consider the possibility of rival apsychologistic approaches.

Is there an acceptable view of a priori knowledge which rivals the Kantian conception? The logical positivists hoped to understand a priori knowledge without dabbling in psychology. The simplest of their proposals was the suggestion that X knows a priori that p if and only if X believes that p and p is analytically true.<sup>27</sup>

Gilbert Harman has argued cogently that, in cases of factual belief, the nature of the reasons for which a person believes is relevant to the question of whether he has knowledge.<sup>28</sup> Similar considerations arise with respect to propositions which the positivists took to be a priori. Analytic propositions, like synthetic propositions, can be believed for bad reasons, or for no reasons at all, and, when this occurs, we should deny that the believer knows the propositions in question. Assume, as the positivists did, that mathematics is analytic, and imagine a mathematician who comes to believe that some unobvious theorem is true. This belief is exhibited in her continued efforts to prove the theorem. Finally, she succeeds. We naturally describe her progress by saying that she has come to know something she only believed before. The positivistic proposal forces us to attribute knowledge from the beginning. Worse still, we can imagine that the mathematician has many colleagues who believe the theorem because of dreams, trances, fits of Pythagorean ecstasy, and so forth. Not only does the positivistic approach fail to separate the mathematician after she has found the proof from her younger self, but it also gives her the same status as her colleagues.

A natural modification suggests itself: distinguish among the class of analytic truths those which are elementary (basic laws of logic, immediate consequences of definitions, and, perhaps, a few others), and propose that elementary analytic truths can be known merely by being believed, while the

rest are known, when they are known a priori, by inference from such truths. Even this restricted version of the original claim is vulnerable. If you believe the basic laws of logic because you have learned them from an eminent mathematician who has deluded himself into believing that the system of *Grundgesetze* is consistent and true, then you do not have a priori knowledge of those laws. Your belief in the laws of logic is undermined by evidence which you do not currently possess, namely the evidence which would expose your teacher as a misguided fanatic. The moral is obvious: apychologistic approaches to a priori knowledge fail because, for a priori knowledge as for factual knowledge, the reasons for which a person believes are relevant to the question of whether he knows.

Although horror of psychologizing prevented the positivists from offering a defensible account of a priori knowledge, I think that my analysis can be used to articulate most of the doctrines that they wished to defend. Indeed, I believe that many classical theses, arguments and debates can be illuminated by applying the analysis presented here. My aim has been to prepare the way for investigations of traditional claims and disputes by developing in some detail Kant's conception of a priori knowledge. "A priori" has too often been a label which philosophers could attach to propositions they favored, without any clear criterion for doing so. I hope to have shown how a more systematic practice is possible.

## NOTES

I am grateful to several members of the Department of Philosophy at the University of Michigan for their helpful comments on a previous version of this paper, and especially to Alvin Goldman and Jaegwon Kim for their constructive suggestions. I would also like to thank Paul Benacerraf, who first interested me in the problem of characterizing a priori knowledge, and prevented many errors in early analyses. Above all, I am indebted to Patricia Kitcher and George Sher, who have helped me to clarify my ideas on this topic. Patricia Kitcher's advice on issues in the philosophy of mind relevant to §V was particularly valuable.

1. There are some exceptions. Passing attention to the problem of defining apriority is given in Pollock 1974, chapter 10; Swinburne 1975, esp. pp. 238–241; Erwin 1974, especially pp. 593–597. The inadequacy of much traditional thinking about apriority is forcefully presented in Kripke 1971, especially pp. 149–151, and Kripke 1972, especially pp. 260–264.

2. See Kripke 1971, 1972.

3. For ease of reference, I take propositions to be the objects of belief and knowledge, and to be what declarative sentences express. I trust that my conclusions would survive any successful elimination of propositions in favor of some alternative approach to the objects of belief and knowledge.

4. Kant 1787, B2–3.

5. See Pollock 1974; Swinburne 1975; Erwin 1974.

6. Someone might be tempted to propose, conversely, that all innate knowledge is a priori (cf. Swinburne 1975, p. 239). In Kitcher 1978, I have argued that there may well be no innate knowledge and that, if there were any such knowledge, it would not have to be a priori.

7. Pollock 1974, p. 301, claims that we can only resist the suggestion that this knowledge is independent of experience by complicating the notion of experience. For the reason given in the text, such desperate measures seem to me to be unnecessary.

8. See the example of the man who undermines the foundations of his house, Kant 1787, B3.

9. Prominent exponents of this approach are Alvin Goldman, Gilbert Harman and David Armstrong. See Goldman 1967, 1975, 1976, 1979; Harman 1973; Armstrong 1973.

10. Psychologistic epistemologies are often accused of confusing the context of discovery with the context of justification. For a recent formulation of this type of objection, see Lehrer 1974, pp. 123 ff. I have tried to show that psychologistic epistemology is not committed to mistakes with which it is frequently associated in Kitcher 1979. I shall consider the possibility of an apychologistic approach to apriority in §VII below.

11. More details about Kant's theory of pure intuition can be found in my paper Kitcher 1975, especially pp. 28–33.

12. For further discussion of this requirement and the possibility of the contingent a priori, see §V below.

13. For different reasons, Goldman proposes that an analysis of the general notion of warrant (or, in his terms, justification) can focus on psychological processes. See section 2 of Goldman 1979.

14. Consider, for example, a Kantian process of pure intuition which begins with the construction of a triangle. Should we say that a process of the same type must begin with the construction of a triangle of the same size and shape, a triangle of the same shape, any triangle, or something even more general? Obviously there are many natural classifications here, and I think the best strategy is to suppose that an apriorist is entitled to pick any of them. Strictly, the sets which do not constitute types are those which violate correct taxonomies. In making present decisions about types, we

assume that our current principles of classification are correct. If it should turn out that those principles require revision then our judgments about types will have to be revised accordingly.

16. For a lucid and entertaining presentation of this point, see Lewis 1976, pp. 149–151.

17. Of course, X might have been more intelligent, that is, he might have had better versions of the faculties he has. We allow for this type of change. But we are not interested in worlds where X has extra faculties.

18. The idea that "actual" is indexical is defended by David Lewis in Lewis 1970. In van Fraassen 1977, Bas van Fraassen puts Lewis' ideas about "actual" in a general context. The machinery which van Fraassen presents in that paper can be used to elaborate the ideas of the present paragraph.

19. Jaegwon Kim has pointed out to me that, besides the "species-relative" notion of apriority presented in the text, there might be an absolute notion. Perhaps there is a class of propositions which would be knowable a priori by any being whom we would count as a rational being. Absolute a priori knowledge would thus be that a priori knowledge which is available to all possible knowers.

20. In characterizing pain as an involuntary state one paragraph back I may seem to have underestimated our powers of self-torture. But even a masochist could be defeated by unkind experience: as he goes to pinch himself his skin is anesthetized.

21. I shall ignore the tricky issue of trying to say exactly what is known when we know this and kindred things. For interesting explorations of this area, see Casteneda 1967, 1968; Perry 1977, 1979. The issue of how to represent the content of items of self-knowledge may force revision of the position taken in footnote 3 above: it may not be possible to identify objects of belief with meanings of sentences. Although such a revision would complicate my analysis, I don't think it would necessitate any fundamental modifications.

22. This presupposes that our knowledge of our existence does not result from some special kind of "inner sensation." For, if it did, different lives would deprive us of the warrant.

23. Kripke 1971, 1972 has attempted to construct examples of contingent propositions which can be known a priori. I have not tried to decide here whether his examples are successful, since full treatment of this question would lead into issues about the analysis of the propositions in question which are well beyond the scope of the present paper. For a discussion of some of the difficulties involved in Kripke's examples, see Donnellan 1977.

24. Kripke seems to take this to be the main argument against the contingent a priori. See Kripke 1972, p. 263.

25. As the discussion of this paragraph suggests, there is an intimate relation between my requirements (2(b)) and (2(c)). Indeed, one might argue that (2(b)) would not be met unless (2(c)) were also satisfied—on the grounds that one cannot allow a process to override experience unless it guarantees truth. The subsequent

discussion will show that this type of reasoning is more complicated than appears. Hence, although I believe that the idea that a priori warrants function independently of experience does have implications for the reliability of these processes, I have chosen to add (2(c)) as a separate condition.

26. This notion of contextual apriority has been used by Hilary Putnam. See, for example, Putnam 1975a, 1978, especially p. 154.

27. See Ayer 1936, chapter IV and Schlick 1959, especially p. 224.

28. Harman 1973, chapter 2; see also Goldman 1979, section 1.

## BIBLIOGRAPHY

- ARMSTRONG, DAVID M., 1973, *Belief, Truth and Knowledge*. London: Cambridge University Press.
- AYER, A. J., 1936, *Language Truth and Logic*. London: Victor Gollancz. Ltd.
- CASTENEDA, HECTOR-NERI, 1967, "Indicators and Quasi-indicators," *American Philosophical Quarterly* 4, 85–100.
- 1968, "On the Logic of Attributions of Self-Knowledge to Others," *Journal of Philosophy* 64, 439–56.
- DONNELLAN, KEITH, 1977, "The Contingent A Priori and Rigid Designators," in *Contemporary Perspectives in the Philosophy of Language*, ed. P. French, T. Uehling, Jr., and H. Wettstein. Minneapolis Minn.: University of Minnesota Press.
- ERWIN, E., 1974, "Are the Notions 'A Priori Truth' and 'Necessary Truth' Extensionally Equivalent?" *Canadian Journal of Philosophy* 3, 591–602.
- GOLDMAN, ALVIN I., 1967, "A Causal Theory of Knowing," *Journal of Philosophy* 64, 357–372. (Reprinted in this volume.)
- 1975, "Innate Knowledge," in Stich 1975.
- 1976, "Discrimination and Perceptual Knowledge," in *Journal of Philosophy*, 73, Nov 1976, 771–791.
- 1979, "Varieties of Cognitive Appraisal," *Nous* 13, 23–38.
- HARMAN, GILBERT, 1973, *Thought*. Princeton, N.J.: Princeton University Press.
- KANT, IMMANUEL, 1787, *Critique of Pure Reason*.

- KITCHER, PHILIP, 1975, "Kant and the Foundations of Mathematics," *The Philosophical Review* 84, 23–50.
- 1978, "The Nativist's Dilemma," *Philosophical Quarterly* 28, 1–16.
- 1979, "Frieze's Epistemology" in *Philosophical Review*, 88, April 1979, 235–262.
- KRIPKE, SAUL, 1971, "Identity and Necessity," in *Identity and Individuation*, ed. M. Munitz. New York: New York University Press.
- 1972, "Naming and Necessity," in *Semantics of Natural Language*, ed. D. Davidson and G. Harman. (Dordrecht, the Netherlands: D. Reidel Publishing Co.)
- LEHRER, KEITH, 1974, *Knowledge*. Oxford: Oxford University Press.
- LEWIS, D., 1970, "Anselm and Actuality," *Nous* 4, 175–88.
- 1976, "The Paradoxes of Time Travel," *American Philosophical Quarterly* 13, 145–52.
- PERRY, J., 1977, "Frege on Demonstratives," *Philosophical Review* 86, 474–97.
- 1979, "The Problem of the Essential Indexical," *Nous* 13, 3–21.
- POLLOCK, JOHN, 1974, *Knowledge and Justification*. Princeton, N.J.: Princeton University Press.
- PUTNAM, HILARY, 1975, *Mathematics, Matter and Method*. Cambridge University Press.
- SCHLICK, MORITZ, 1959, "The Foundation of Knowing," in *Logical Positivism*, ed. A. J. Ayer. New York: The Free Press.
- STICH, STEPHEN, *Innate Ideas*, 1975, Berkeley, CA: University of California Press.
- SWINBURNE, R. G., 1975, "Analyticity, Necessity and Apriority," *Mind* 84, 225–43.
- VAN FRAASSEN, Bas, 1977, "The Only Necessity Is Verbal Necessity," *Journal of Philosophy* 74, 71–85.
- BERTOLET, ROD, "Putnam on the A Priori," *Philosophia*, 18 (July 1988).
- CASULLO, ALBERT, "The Definition of A Priori Knowledge," *Philosophy and Phenomenological Research* 38 (1977).
- CHOMSKY, N., H. PUTNAM, and N. GOODMAN, "Symposium on Innate Ideas." In *Boston Studies in the Philosophy of Science*, Vol. 3, ed. R. S. Cohen and M. W. Wartofsky. Dordrecht; The Netherlands: D. Reidel Publishing Co., 1967.
- FREGE, GOTTLLOB, "Are the Laws of Arithmetic Synthetic A Priori or Analytic?" In *The Foundations of Arithmetic*, 2nd edition, Gottlob Frege. New York: Harper & Row, Publishers, Inc., 1953.
- GASKING, DOUGLAS, "The Analytic–Synthetic Controversy," *Australasian Journal of Philosophy* 50 (1972).
- GRICE, H. P., and P. F. STRAWSON, "In Defense of a Dogma," *Philosophical Review* 65 (1956).
- HAMLYN, D. W., "A Priori and A Posteriori." In *The Encyclopedia of Philosophy*, Vol. 1, ed. Paul Edwards. New York: Macmillan Publishing Company, 1967.
- , "Analytic and Synthetic Statements." In *The Encyclopedia of Philosophy*, Vol. 1, ed. Paul Edwards. New York: Macmillan Publishing Company, 1967.
- , "Contingent and Necessary Statements." In *The Encyclopedia of Philosophy*, Vol. 2, ed. Paul Edwards. New York: Macmillan Publishing Company, 1967.
- HARRIS JAMES F. JR., and RICHARD H. SEVERENS, ed. *Analyticity*. Chicago: Quadrangle Books, 1970.
- HINTIKKA, JAAKKO, "Are Logical Truths Analytic?" *The Philosophical Review* 74 (1965).
- KENNEDY, RALPH, "Salmon versus Kripke on the A Priori," *Analysis* 47 (June 1987).
- KIM, JAE GWON, "The Role of Perception in 'A Priori' Knowledge: Some Remarks," *Philosophical Studies* 40 (1981).
- KITCHER, PHILIP, "Apriority and Necessity," *Australasian Journal of Philosophy* 58 (1980).
- KRIPKE, SAUL, "Naming and Necessity." In *Semantics of Natural Languages*, ed. D. Davidson and G. Harman. Dordrecht, The Netherlands: D. Reidel Publishing Co., 1972.
- LANGFORD, C. H., "A Proof That Synthetic A Priori Propositions Exist," *Journal of Philosophy* 46 (1949).

## SELECTED BIBLIOGRAPHY

- AYER, A. J., "Basic Propositions," in *Philosophical Analysis*, ed. Max Black. Ithaca, N.Y.: Cornell University Press, 1950.
- BENACERRAF, PAUL, "Mathematical Truth," *Journal of Philosophy*, 70 (1973).
- BENNETT, JONATHAN, "Analytic–Synthetic," *Proceedings of the Aristotelian Society* 59 (1959).

- MACKIE, J. L., "The Possibility of Innate Knowledge," *Proceedings of the Aristotelian Society* 70 (1969-70).
- MOSER, PAUL K., ed., *A Priori Knowledge*. New York: Oxford University Press, 1987.
- PAP, ARTHUR, *Semantics and Necessary Truth*. New Haven, Conn.: Yale University Press, 1958.
- QUINE, W. V., "Truth by Convention." In *The Ways of Paradox*, revised. Cambridge, Mass.: Harvard University Press, 1976.
- ROSENTHAL, SANDRA, *The Pragmatic A Priori: A Study in the Epistemology of C. I. Lewis*. St. Louis, Mo.: Warren H. Green Inc. 1976.
- SLEIGH, R. C., ed. *Necessary Truth*. Englewood Cliffs, N.J.: Prentice Hall, 1972.
- STICH, S., ed. *Innate Ideas*. Berkeley Calif.: University of California Press, 1975.
- WANG, HAO, "Two Commandments of Analytic Empiricism," *Journal of Philosophy* LXXXII, 9 (September 1985).
- WHITE, MORTON, "The Analytic and the Synthetic: An Untenable Dualism." In *Semantics and the Philosophy of Language*, ed. Leonard Linsky. Urbana, Ill.: University of Illinois Press, 1952.