

A PRIORI KNOWLEDGE REVISITED

Philip Kitcher

I

Some years ago, I offered an account of a priori knowledge. My aim in doing so was to prepare for inquiring if mathematical knowledge is, or can be, a priori. Since I ended up defending an unpopular answer to this question—"No"—it's hardly surprising that people have scrutinized the account, or that many have concluded that I stacked the deck in the first place. Of course, this was not my view of the matter. My own judgment was that I'd uncovered the tacit commitments of mathematical apriorists and that the widespread acceptance of mathematical apriorism rested on failure to ask what was needed for knowledge to be a priori. Nevertheless, my critics have raised important challenges, and have offered rival conceptions that are less demanding. I want to examine their objections to my explication of a priori knowledge, and to explore whether the weaker alternatives succeed in preserving traditional philosophical claims. What follows is a mixture of penitence and intransigence.

Let me start by reviewing some features of my account and adding connections with other important notions, only some of which I've noted before. First, I take 'a priori' to be an epistemological predicate, one that applies primarily to items of knowledge, so that the notion of a priori truth is derivative: a priori truths are those propositions that could be known a priori. It seems to me that analyses that begin with the conception of a priori truth run the risk of conflating apriority with something quite different – the property of being a logical truth, or a proposition whose truth can be explained by features of concepts, for example.

The account of a priori knowledge is embedded within a general approach to knowledge (the psychologistic approach) according to which whether or not a state of true belief counts as a state of knowledge depends on the causal process that generated that state. If a state is produced by the right kind of causal process, so that it is a state of knowledge, then I say that the process is a warrant for the belief. My general understanding of warrants is a version of reliabilism (of which more later): warrants are processes belonging to types that regularly and reliably produce true belief. One condition that I don't require of a warrant is that it involve reasons that the knower could cite in justification of the belief. This idea would prove important were we to consider if denying the apriority of logic and mathematics commits one to skepticism.

A priori knowledge is knowledge produced by special types of processes, a priori warrants. The classical way of introducing the notion of a priori knowledge is to take it to be knowledge that is independent of experience. However, most apriorists have allowed for the possibility that we could have a priori knowledge of propositions containing concepts that could only be garnered from experience. Hence I explicate a priori knowledge by reference to the notion of an experience's being sufficiently rich for p, that is that someone who had that experience could acquire the concepts needed to entertain p. Now we can't say that X knows a priori that p just in case X knows that p, and, for any e sufficiently rich for p, X could have had e and known that p, because this formulation would collapse the distinction between a priori knowledge and empirical knowledge of propositions that could have been known a priori. To arrive at an appropriate definition, we need to insist that the subject could have known the proposition in the same way, given any sufficiently rich experience, and this leads us to the idea of an a priori warrant, a process that could have produced the knowledge against the background of any sufficiently rich experience.

So, finally, the account:

X knows a priori that p iff. X knows that p and X's knowledge that p was produced by a process a that is an a priori warrant for p.

a is an a priori warrant for X's belief that p just in case a is a process such that for any sequence of experiences sufficiently rich for X for p

(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.

The explication of a priori knowledge in terms of a priori warrants is only likely to be controversial among those who object to the background psychologistic framework. I won't repeat the arguments that have been offered in support of this framework; instead, I'll focus on criticisms of my conditions on a priori warrants.

Of these, (a) is the least controversial. The motivation for (a) draws on the traditional idea that a priori knowledge must be independent of experience: surely, if the process that undergirds one's knowledge would not have been available, given a different sufficiently rich experience, then that knowledge can't count as independent of experience. Nonetheless there's an important objection. After all, might it not be the case that, in some worlds, while I'm able to grasp the concepts pertinent to a proposition, I lack the intellectual acumen to formulate the proof that I actually construct; or perhaps my memory lets me down, and I don't discover a complicated, tree-structured, justification? It seems plausible to suppose that these possibilities don't undermine the status of the proof as an a priori warrant in the actual world. This line of criticism reveals that we must be quite careful in interpreting (a).

To assess the credentials of an actual process a as an a priori warrant for my belief that p , we're asked to consider possible worlds in which I have different experiences that are sufficiently rich for me for p . Let e be one such sequence of experiences. Then there's a set of worlds, $W(e)$, in which e is exactly my sequence of experiences. To say that some process of the same type as a would be available given e is to maintain that there's some world in $W(e)$ in which I produce a process of that type. That's quite compatible with supposing that in many worlds in $W(e)$ I am appallingly dull, or have a memory like a sieve, or am simply never interested enough to seek out the pertinent type of process. Indeed, it seems that $W(e)$ will contain a world very similar to the actual one, in which my faculties are much as they are in the actual world but my actual stream of experience is replaced by e , and in which I undergo a process of the same type as a (construct the same proof). Condition (a) thus allows actual processes to survive the envisaged possibilities.

There is, nonetheless, a deeper point to the criticism. Isn't it possible that the proof of p might require concepts beyond those occurring in p , and that, consequently, some experiences sufficiently rich for me for p might not be sufficiently rich for the formulation of the proof of p ? Specifically, let p contain the concepts in C and suppose that the concepts in $C+$ ($C \dot{\cup} C+$) are required to formulate the proof. Let e now be an experience that suffices for my acquisition of all the concepts in C but not for all those in $C+$. There's now no world in $W(e)$ in which I undergo a process of the same type as the proof. Hence (a) rules out the actual proof as an a priori warrant.

I offer two possible reactions to this point. First, one might declare that this is just as it should be. Behind this declaration is the thought that the classical conception of apriority supposes that experiences that would suffice for entertaining a proposition are also enough to allow a thinker who is astute to generate the conceptual resources for the proof of that proposition – if that can't be done, then the alleged proof isn't an a priori warrant but a process that depends on special features of the person's experience to supply the needed concepts. In case that response seems too blunt, one can accommodate the example by a slight amendment to my account. Instead of supposing that the crucial notion for recognizing the role of experience in a priori knowledge in the supplying of concepts is that of a life sufficient for X for p , where p is the proposition known, we can replace that notion with that of a life sufficient for X for a $*$ (where a $*$ is the type to which a belongs). Obviously this substitution rules out the worrying possibility that some lives wouldn't allow for the formulation of the proof, since, for any sufficiently rich life e , $W(e)$ now contains worlds in which the knower has all the pertinent concepts. I won't try to decide between these alternatives, but will leave it to the friends of the a priori to decide which more adequately captures the venerable idea that experience may be needed for people to have the conceptual resources to gain a priori knowledge.

I conclude that condition (a), either in its original form or slightly amended, is unproblematic. Conditions (b) and (c), on the other hand, have aroused objections that require much more lengthy discussion. Many of my critics have wanted to adopt an account of a priori knowledge that makes a priori warrants

fallible or defeasible. So, in the ensuing sections, I'll want to compare two conceptions of a priori knowledge: the Strong Conception that requires (a)-(c), and the Weak Conception that demands only (a). The strategy will be to examine the conceptions in relation to traditional thinking about apriority, looking both at general theses about links to other philosophical notions and at the particular types of knowledge that are classified as a priori. I'll start with the former.

II

In a classic discussion, Saul Kripke revealed deep problems with the thesis

- (1) It is possible to know a priori that p iff. it's necessary that p .

Given either the Weak or the Strong conception, (1) fails in both directions. For there are plenty of necessary truths that we are unable to express, let alone know a priori; conversely, each of us can know a priori that he/she exists and that he/she is actual; the latter knowledge can be used to generate Kripke-style examples of contingent a priori knowledge.

Next, consider

- (2) It is possible to know a priori that p iff. p is known innately.

On the Strong conception, (2) fails in both directions. First, note that the relativization to lives sufficiently rich for p already allows for a priori knowledge that isn't innate; conversely, even though one could know that p on the basis of no experience, it doesn't follow that one could know that p on the basis of any sufficiently rich experience; we might, for example, have innate knowledge of some principle of universal grammar, but later acquire empirical evidence for the existence of a language spoken by a fellow human (or humanoid) that violated that principle; whatever process originally warranted our belief would not be able to discharge its warranting function against this experiential background (condition (b) would fail). The Weak conception rejects (2) on the grounds that not all a priori knowledge is innate, but it does preserve the traditional idea that innate knowledge is a priori. On this score, then, the Weak conception seems closer to the tradition.

Matters are similar with respect to

- (3) It is possible to know a priori that p just in case p is analytic.

The Strong conception takes this to fail in both directions. On the one hand, it allows that "I exist" can be known a priori by each of us, but in none of our mouths is this analytic; on the other, if we think of analytic truths as those which can be known on the basis of our grasp of pertinent concepts, there is an old point, traceable to Mill and Kant, that experience can call into question our right to employ concepts and thus undermine the warranting power of a

process that stems from our grasp of concepts; consequently processes that elaborate our grasp of concepts don't meet condition (b) on a priori warrants in all instances (perhaps not even in many). The Weak conception agrees with the Strong conception about the first type of failure of (3), concluding that not all propositions that are knowable a priori are analytic, but, because it allows for defeasible a priori warrants and rejects my condition (b), it treats the Mill-Kant point differently, supposing that, when experience doesn't undermine the use of the pertinent concepts, our grasp of those concepts can serve to generate an a priori warrant for analytic truths. Thus, the Weak conception retains one of the most central ideas of traditional thinking about apriority, the thesis that grasp of concepts will typically yield a priori knowledge.

The Strong conception can manage, at best, a restricted version of this thesis, for propositions whose constituent concepts could never be undermined by experience. With a nod to Kant, let's call these a priori concepts. An obvious explication is

C is an a priori concept just in case, for any sufficiently rich experience e, X would be warranted in believing that C was apt for the description of experience.

(Here I assume a form of epistemic democracy – what is an a priori concept for one is an a priori concept for all.) Now it does follow that if there are a priori concepts, then analytic propositions containing only such concepts (provided there are any) can be known a priori. It seems to me that one of Kant's great insights was the point that analyticity isn't a solution to the problem of a priori knowledge except in very special cases.

However, Kant (and some of his successors) have often conflated the notion of an a priori concept with a related, but distinct, idea. Let's say that a conceptual precondition for experience (or knowledge) is a concept that a person must use in judgments if that person is to have experience (or knowledge). Assuming epistemic democracy, again, we can inquire if the Strong conception supports the thesis

(4) C is a conceptual precondition of experience (knowledge) just in case C is an a priori concept.

It doesn't. For, even though we might have to deploy a concept in order to have experience (knowledge), it doesn't follow that our belief that that concept was apt for the description of experience would have to be warranted against the background of any sufficiently rich experience. There are surely experiences that could dupe us into thinking that this is wrong, since finding the conceptual preconditions of experience (knowledge) has been a hotly contested business in the history of philosophy. Conversely, even though we could know that a concept was apt for the description of experience against the background of any sufficiently rich experience, it would not follow that that concept was a precondition of experience (knowledge). Perhaps recondite mathematical concepts are such that we would always be warranted in believing of them that they

are apt for the description of experience, but we don't have to employ them to have experience (knowledge). Indeed, this may even be true of the simplest mathematical concepts.

In similar fashion, we can consider propositional preconditions of experience where

p is a propositional precondition of experience (knowledge) just in case we have to believe that p in order to have experience (knowledge).

The Strong conception rejects

(5) p is a propositional precondition of experience (knowledge) just in case p can be known a priori.

If there are any propositions that we have to believe in order to have experience (knowledge), it's an entirely separate issue whether there are processes that would warrant them given any sufficiently rich experience. This means that there are experiences, given which we'd be warranted in not believing propositions that are required for us to have experience (knowledge), and that, of course, sounds rather dire. As we shall see, however, this doesn't mean that we couldn't continue to have something rather like experience (knowledge). The failure of (5) in the opposite direction is more straightforward, since there are surely pieces of a priori knowledge that we don't have to believe in order to have experience – to use an earlier example, each of us can know a priori that he/she is actual, but I find it highly dubious that we have to believe this to have experience (knowledge).

The score so far is heavily in favor of the Weak conception. Not only does it preserve the classical ideas that innate knowledge and knowledge based on grasp of concepts are both a priori, but the Strong conception appears at odds with influential attempts to preserve residues of the latter theme. Hence, at least at first sight, the Strong conception appears vulnerable to complaints that it stacks the deck against the a priori.

III

So what can be said in its favor? When I originally developed my account of a priori knowledge, I claimed that conditions (b) and (c) were an important part of the idea that a priori knowledge is independent of experience. In fact, it seems to me that we can simplify the task of appraising the Strong conception by focusing just on (b), since, if (b) were accepted, it would be hard to resist (c).

Here's why. Suppose that there were processes satisfying (a) and (b) but not (c). Then there could be experiences sufficient for p such that, given those experiences, one might be warranted in believing that p, even though p were false. Let e be any such experience, and consider the following extension of e, e*:

the next phase of one's life consists in an encounter with an oracle who demonstrates power to answer vast numbers of significant questions, who testifies to the falsity of p , and who offers whatever can be done to show directly that p is false. e^* is sufficient for p , so, by assumption, the process is available given e^* and its warranting power is unaffected. But, in order to believe that p , one must override extremely strong evidence to the contrary. How can the process give one a licence to override when one would have been epistemically better off not overriding? Intuitively, the person who overrides the performance of the oracle is being dogmatic in insisting on the belief generated by the process, and the belief is no longer warranted. Hence it seems to me that the process cannot satisfy (b) unless it also satisfies (c).

One might try to avoid this argument by finding a proposition p and a belief-generating process such that the falsity of p and the existence of the process are incompatible with the assumptions used in the extension of e to e^* . But I have not been able to think of a convincing case in which this is so. I conclude that (c) is, in fact, a consequence of my other conditions on a priori knowledge. Thus, if there were a cogent defense of (b), it would provide a rationale for the Strong conception.

I used to believe that there was a relatively simple argument that would do the trick. After all, the heart of the classical view of apriority is that a priori knowledge is knowledge independent of experience. But it seems that one can hardly claim that an a priori warrant is independent of experience if its production of true belief and its warranting power could be subverted in worlds where one had different experiences. In various places, I tried to support this idea along the following lines:

... if alternative experiences could undermine one's knowledge then there are features of one's current experience which are relevant to the knowledge, namely those features whose absence would change the current experience into the subversive experience. The idea of the support lent by kindly experience is the obverse of the idea of the defeat brought by uncooperative experience.

But this argument is bad. As several people have pointed out, there are important, epistemologically relevant, differences between the roles played by subversive experiences and those played in normal cases by the absence of those conditions. To take an example of Charles Parsons', the fact that under imaginable circumstances I could have grounds for believing that my perceptual experience isn't veridical doesn't entail that the absence of those experiences are now playing a causal role in generating or sustaining my perceptual beliefs – I don't have to convince myself that I'm not the dupe of a Cartesian demon every time I observe objects in my vicinity.

Indeed (I blush to admit) nobody who advocates a reliabilist approach to knowledge ought to be at all tempted by the reasoning just rehearsed. Without going very far into the details of the position, we can characterize reliabilism as the claim that the processes that generate and sustain those beliefs that count as knowledge are processes that regularly (though not invariably) generate true

belief in the class of relevant alternatives to the actual situation. In introducing the position, Alvin Goldman explicitly considered ordinary perception in a variety of contexts. On most occasions, when people drive through the countryside and observe a barn, they come to know that there's a barn in the vicinity. This is because the process that generates the belief that there's a barn over there belongs to a type that regularly generates true beliefs across relevant alternatives. If, however, this particular region is used as a movie set (perhaps for films based on Hardy's novels) and is full of barn-facsimiles, then, even if the particular object in the line of vision is a genuine barn, the tourist doesn't know that it's a barn, precisely because, in this instance, the class of relevant alternatives includes those occasions in which his eyes were pointed towards barn facsimiles. The crucial point here is that to appraise a candidate process as a warrant one must consider whether processes of the type would regularly generate true beliefs across the class of relevant alternatives, and that class varies according to facts about the world (as in the presence of barn facsimiles) and according to facts about the person (as when someone has, incorrectly, told him that this region is full of barn facsimiles). This has direct significance for claims about a priori warrants. Provided that background conditions are benign, the class of alternatives is restricted and a type of process might well succeed in regularly generating true beliefs across this class, even though, were experience to be recalcitrant, a wider class of alternatives would be relevant, across which that type of process would fail the test of regularly generating true belief. Reliabilism thus develops the point that Parsons lucidly made against the Strong conception, and supports the thought that a process available given any sufficiently rich life might warrant belief under actual circumstances even though its warranting power could be defeated by some possible streams of experience, the thought that is central to the Weak conception.

Indeed, the idea of defeasible a priori warrants for mathematical knowledge can be developed quite plausibly on the basis of reliabilism. Consider two different versions of the thesis that mathematical knowledge is a priori. The first concentrates on the knowledge of new mathematical principles by the great mathematicians who discover (introduce?) them – we envisage Gauss or Dedekind or Cantor coming to a priori knowledge that nobody has had before on the basis of some kind of process (call it 'intuition'). In response to the challenge that there are circumstances under which the mathematician in question might have had reasons to worry about the reliability of the intuition, reliabilism notes that, under the actual circumstances, defeating such doubts may not be relevant, so that the intuition serves as an a priori warrant. The second version focuses on the elementary arithmetical knowledge possessed by everyone, claiming that once the basic concepts of arithmetic have been acquired each of us knows such things as that $2 + 1 = 3$ on the basis of our grasp of the pertinent concepts. Faced with the Mill-Kant point that experience might reveal that the employment of a concept (or family of concepts) isn't warranted, reliabilism concedes that, given certain kinds of experiences, the process of warrant by grasp of concepts would be defeated, but maintains that, for all of us virtually all of the time, there's no reason to worry about the concepts we've acquired,

so that alternatives in which we face experiential challenges to our conceptual framework aren't relevant, and the process thus succeeds in warranting belief under actual circumstances.

I conclude that my earlier defense of the Strong conception and the attempts to refute the apriority of mathematics both fail. Nonetheless, I think that matters are more complicated than they've so far appeared. The Weak conception has a somewhat problematic relation to classical views of the a priori, and it will be important to draw out some of its epistemological consequences. In the next section, I'll take up some preliminary points. Sections V and VI will take up further issues that I view as more fundamental.

IV

I'll start with one of the substantive views of a priori mathematical knowledge envisaged in the last section. We imagine people like Cantor or Dedekind or Gauss introducing new principles on the basis of some special apprehension of mathematical truth (of course, there's a well-known puzzle of Paul Benacerraf's about how this could possibly work, but that's not my main concern here). Now the most obvious difference between this and ordinary perception is that the ability is very rare indeed. Further, because it's rare, it's very difficult to get other people to check on whether or not you have done it well. Lastly, the exercise of this alleged ability has a highly chequered history – it's not just the Frege-Cantor disaster with naive set theory that ought to concern us, but the long sequence of mathematicians who announced principles about continuity, limits, infinity and similar notions, whose ideas were, more or less quickly, repudiated. Worries about ordinary empirical knowledge arise only in special circumstances, but that's because each of us has ample background reinforcement from others. By contrast, appeals to elusive processes of a priori reason ought always to be accompanied by doubts about whether one has carried out the process correctly, and whether, in this instance, the deliverances are true. Hence it appears that when we consider the kinds of processes that might have given a priori warrants for major innovations in mathematics the class of relevant alternatives over which those processes should regularly produce true belief is considerably expanded, including many instances in which the mathematician engages in beguiling procedures that lead to error. The power to warrant belief in the actual situation would thus be undermined – and indeed, we might claim that one couldn't satisfy the Weak conception unless the Strong conception were also satisfied.

My critics could reasonably reply that I've stacked the deck by picking out the least plausible candidates for a priori warrants – Gödelian intuitions and kindred processes. The second substantive apriorist conception of the last section, they may insist, is free of the troubles I've canvassed. We acquire the basic arithmetical concepts at teacher's knee, and from then on carry within us the ability to apprehend such things as that $2 + 1 = 3$. Unlike the cloudy deliverances of mathematical intuition, this mundane grounding of arithmetical

knowledge in our grasp of concepts has an auspicious track record – as good, in fact, as that of ordinary perceptual experience. In consequence, there’s no threatening expansion of the class of relevant alternatives and no intensification of the demand that the warranting process be reliable.

But what are we to say about the innovators who fashioned the framework of concepts on which later a priori knowledge is to draw? It seems to me that there are three possibilities. First, one can propose that the innovators’ knowledge was grounded in some a priori intuition, precisely the view I’ve been trying to undermine. Second, in line with the view that a priori knowledge stems from grasp of concepts, it’s possible to claim that the process of fashioning the concepts automatically yields a priori knowledge. That suggestion is vulnerable to an argument exactly parallel to the critique of intuition: the history of mathematics is full of unfruitful, even incoherent, specifications of mathematical concepts, so that, for the aspiring innovator, it’s always relevant to determine how the novel definition serves the purposes of the pertinent branch of mathematics. As Dedekind (to cite just one example) saw so clearly, his definition of continuity had to be supported by a detailed exploration of its consequences. Once again, the expansion of the class of relevant alternatives undermines the claim that the pertinent process – in this instance, conceptual innovation – provides a priori warrant, even on the Weak conception. The third, and most plausible, possibility is to declare that the knowledge of the innovators isn’t a priori: they are warranted in proposing new concepts and principles through an often lengthy process of demonstrating that their new ideas play a fruitful role within inquiry. But now we arrive at a curious picture: a particular framework of concepts is available to us because of the history of constructing those concepts and showing the experiential fruitfulness of the framework; those who played crucial roles in this history did not have a priori knowledge, but, we, their lucky successors, are able to do better. The Weak conception claims that it doesn’t matter that our current processes of generating belief are dependent on their endeavors. Note, however, that we might reasonably view our – allegedly a priori – knowledge as dependent on experience, the experience that warranted adoption of the conceptual framework on which we now draw.

This is only a first look at a crucial issue that will occupy us in Section VI, the question of tradition-dependence. I want to close the present section by taking up the second point advertised above, a point that charges the Weak conception with reducing an apparently important distinction to insignificance.

There are lots of distinctions that could be drawn among various kinds of knowledge, based upon the character of the processes that generated the pertinent beliefs. We could – and we do – distinguish knowledge obtained through testimony from knowledge we acquire for ourselves, knowledge based on current perception from knowledge through memory, knowledge deriving from the use of our eyes from knowledge obtained by using our ears. Dividing items of knowledge according to the causal factors that generated the knowledge might be useful for a number of philosophical and psychological purposes, but none of the divisions I’ve mentioned attributes any special role to a kind of knowledge, a unique function that it can fulfil in our activities as knowers. Historically, the a

priori/a posteriori distinction has had functional significance in that items of a priori knowledge aren't just supposed to be generated in a special way but, once they are in place, to be deployed differently from their empirical counterparts. Indeed, for philosophers like Kant and Frege, the ways in which a priori knowledge is generated are of interest precisely because of the functional significance of the a priori.

The guiding thought is, of course, that a priori principles can be taken for granted in our future empirical investigations. Our scientific hypotheses are revisable – we have to be on the lookout for possible experiences that call them into question – but with respect to our a priori knowledge we don't have to harbor such worries. Hence the a priori/a posteriori distinction has seemed much more fundamental than the distinction between knowledge based on memory and knowledge not so based, or that between knowledge based on touch and knowledge based on smell. Notice, however, that the Weak conception would effectively make the a priori/a posteriori distinction just like these: we'd have another separation of items of knowledge produced by fallible processes, according to the particular faculties that were the source of those processes. We would have abandoned the traditional thought that a priori knowledge can prescribe to experience, that when we know something a priori we don't have to be concerned about what future experiences may bring. Here then is a short argument for the conformity of the Strong conception to the philosophical tradition: the tradition ascribes to a priori knowledge the functional significance of being in a position to prescribe to future experience; knowledge that prescribes to future experience is irrefutable by future experience; any warrant whose warranting power would survive any possible future experience must be a process that would be able to warrant belief given any life sufficiently rich for the pertinent proposition; hence any analysis of apriority that endorses the traditional function of a priori knowledge must satisfy condition (b).

I've been suggesting that the Weak conception is not entirely consonant with traditional thinking about apriority because it abandons parts of the idea that a priori knowledge is independent of experience. The next section will reinforce the complaint by arguing that the Weak conception radically enlarges the scope of our a priori knowledge.

V

Many contemporary apriorists want to use the Weak conception to support the possibility that we can have a priori knowledge generated from our grasp of concepts. But they need to be on guard that the Weak conception doesn't bring them more than they had bargained for. For it seems that all kinds of procedures routinely employed in educational exercises that motivate students to accept various kinds of scientific claims could be transformed into processes that generate a priori knowledge. Let's start with a very simple example.

Imagine a cube on whose faces are inscribed signs for the numbers 1-6, and suppose that this cube is made of some homogeneous material. The cube is

rolled once. What is the chance that the uppermost face will show a 6? Well, one of the faces will show, and the situation is completely symmetrical. Thus there's no reason why one face should be more likely to show than another. Hence the probability that the uppermost face will show a 6 is $1/6$.

Of course, most philosophers have heard this little piece of reasoning many times, and may well have offered it to students who are beginning to think about probabilities. Does it provide those who undergo it with a priori knowledge? Specifically, does it enable people to know a priori that a cubical die made out of uniform material has a chance of $1/6$ of showing 6?

If you think that the answer is "No", then the Weak conception of a priori knowledge runs into fairly immediate trouble. For the process that generates the belief is available given any sufficiently rich life, and the only question is whether that process warrants belief. In fact, given the way the world works, processes of viewing symmetry as a guide to chance regularly generate true beliefs – most dice made of uniform materials conform to this probability. So if we consider the relevant type of process to be the class of thought-experiments in which one employs considerations of symmetry to conclude that the chance that an n -sided die made of uniform materials will land on any particular face will be $1/n$, then it seems that, in the relevant alternatives, cases in which such processes are used to generate beliefs in worlds very like ours, the processes will regularly yield true beliefs. So the symmetry reasoning meets the reliabilist criterion and it should count as a warrant for belief, and, according to the Weak conception, an a priori warrant.

Champions of the Weak conception might try to resist this result by claiming that the pertinent class of processes, or the class of relevant alternatives, are not as I've characterized them. Maybe they propose that we have to consider symmetry arguments in general and to focus on worlds in which such arguments usually fail. The first point to note is that it will be a dangerous and delicate strategy to conjoin this proposal with the claims about a priori knowledge based on grasp of concepts – for it will be important to preserve the easy-going approach that allows us to go along with the concepts we inherit from our teachers without worrying about their scientific utility in our world (or in other worlds). But even if we concede that something like this can be achieved, perhaps because our accumulated wisdom about the use of symmetry arguments casts doubt on their reliability (a move that depends crucially on treating as relevant symmetry arguments that don't take the form of simple reflections on uniform dice), the Weak conception has to recognize that this feature of our tradition is thoroughly contingent.

Let's envisage a world very like our own, one in which the course of nature runs just the same but in which the history of inquiry and the socialization of inquirers are rather different. In this world, the first few phenomena for which people develop symmetry arguments happen to be ones for which such arguments work well, and, in consequence, young investigators are trained to trust the outcome of such arguments – just in the way our tradition sanctions beliefs generated from grasp of concepts. Even if we can't know a priori that a homogeneous cubical die has a probability of $1/6$ of showing 6, our counterparts

in the other world can, and this reveals something rather counterintuitive about the Weak conception of a priori knowledge: what is knowable a priori depends on the history of inquiry and the conditions of socialization.

But perhaps the champions of the Weak conception have another option, that of supposing we can know a priori the probability of a uniform die's showing 6. The obvious danger is that this will set the Weak conception at variance with the classical view of the bounds of apriority – if it delivers this result, then it cheapens the notion. But, almost as obviously, there are two replies: one that seeks to assimilate what is known to what traditional apriorists have thought we could know a priori, and one that contends that there's more a priori knowledge than has been dreamed of in previous rationalist philosophy.

The former starts by asking what is meant by the claim that the die is made out of uniform material. If our everyday standards for testing uniformity are being employed, then, of course, the claim that the die has probability $1/6$ of showing 6 looks synthetic and the classical boundaries of a priori knowledge seem to be outrun. However, it may be suggested that this is the wrong interpretation of what's going on. In fact, the argument works because the standards for uniformity are inflated, so that, in effect, a cube could only be uniform if there was no process that could introduce asymmetries, or, in other words, a die that didn't have probability $1/6$ of showing 6 would be, ipso facto, nonuniform. Hence, the story goes, what is known a priori is analytic.

This interpretation of how the process of reasoning works strikes me as implausible: after all, the point of the reasoning is to induce people to form beliefs about the probabilities of events involving ordinary objects, objects that they can test for uniformity in familiar ways. The result isn't only counterintuitive but also seems to make no sense of the ways in which scientists actually devise experiments to test claims. Instead of checking on the probabilities generated from particular set-ups, what should really be sought are demonstrations that those set-ups accord with the symmetries and invariances posited in the principles that are known a priori. So, in investigating dice, we ought to be exploring whether there are any features of their composition that can be exploited asymmetrically by physical forces.

The last alternative for the Weak conception is to admit that arguments that appeal to symmetries and invariances do warrant belief in substantial claims about nature, and to claim that all this shows is that the classical view of the scope of apriority was a little too narrow. But now we can envisage a smooth transition from my humble example all the way up to claims about sex-ratios in populations, the invariance of laws in frames of reference, even the existence of particles that fill particular places in a mathematical scheme. In all of these instances, people have appealed to symmetry arguments and other nonempirical motivating suggestions to generate conclusions, and these appeals often figure in classroom "demonstrations" for students of the pertinent fields. Unless the proponent of the Weak conception can draw a boundary between the simple probabilistic cases and the thought-experiments of physicists, biologists and economists, there will be no minor addendum to the classical bounding of apriority but a tremendous explosion of "a priori" knowledge. The celebration of

apriority will outrun the scope of a priori knowledge, even for the most ambitious rationalists.

VI

At this point, we come to the last, and, I think, the deepest reason for concern about the Weak conception. In my discussion of mathematical knowledge, I introduced a conception of knowledge that is at odds with central features of the epistemological tradition, although, at the time, I did not see clearly how radical the break was. On my socio-historical conception of knowledge, the knowledge we have today isn't simply a matter of what we have experienced or thought during the course of our lives, but is dependent on the historical tradition in which we stand and on the social institutions that it has bequeathed to us. Each of us acquires in childhood both a stock of beliefs, the lore that our elders pass on, and a training that prompts us to undergo certain sorts of belief-forming processes and to trust those processes as reliable guides to truth. During our lives we modify those beliefs, typically in idiosyncratic and ephemeral ways, occasionally by introducing or discarding beliefs that belong to the transmitted corpus. In the special case of mathematics, I proposed that our knowledge starts with the vast amount of instruction we are given, by teachers who effectively sum up two millennia of accomplishment. The creative mathematicians of our generation will show how some of the assumptions they have acquired generate, in ways governed by the rules they have learned, new, and maybe surprising, conclusions. Maybe a very few of them will introduce some new basic assumption that will be passed on to the next generation. If so, that assumption will be accepted because it fits well with the mathematics that had already been developed beforehand.

I doubt if any figure in the history of epistemology, from Descartes on, has doubted that something like this is the way in which most people actually acquire and modify their beliefs. Yet, ever since Descartes, philosophers have been beguiled by the idea that this isn't the way knowledge should be. Descartes was quite explicit about the ideal of retreating from the beliefs and practices one had acquired, subjecting them to scrutiny, and building them up anew. Even philosophers who don't subscribe to his standards for the reconstructive work still contemplate the possibility of some sort of synchronic justification of belief. Like Descartes, they contrast the ordinary ways in which we come to believe what we do with an epistemological ideal, a synchronic reconstruction of an individual's body of beliefs that shows how and to what extent that person is justified. My radical break with the tradition consisted in abandoning this ideal, in saying, if you like, that the socio-historical process is all there is.

Let us say that a person's knowledge is independent of socio-historical tradition just in case that person could have had the knowledge, even given socialization in a different tradition, provided only that the socialization made it possible to entertain the propositions known. More formally:

X's knowledge that p is tradition-independent just in case for

any tradition adequate to enable X to entertain p X could have had the sequence of experiences X actually has and known that p.

For many, perhaps almost all, of the alternative traditions envisaged here, X may know what he knows in ways quite different from the ways he actually knows things. The point of insisting on the tradition-independence of knowledge is to suppose that, together, X's experiences and powers of thought would provide warrants for belief that do not depend in any way on the particularities of what X has absorbed from the past. Some versions of the idea would go further than I've done, supposing that the concepts required to entertain the propositions someone knows can usually be gleaned from experience or (perhaps) are innate, so that the clause that restricts traditions to those that suffice for entertaining the pertinent propositions makes no difference and can be omitted.

A thesis shared by most epistemologists from Descartes to the present is that our knowledge is tradition-independent, and one of the primary tasks of epistemology is to reveal the "structure" of knowledge by making that clear. So, one imagines an idealized subject, retreating from society, considering all the propositions she accepts and appealing to a mixture of experience and reason to warrant them for herself. This has been most evident in discussions of empirical knowledge, where, in their different ways, foundationalists and coherence theorists have offered synchronic reconstructions of a subject's beliefs. Now I think that all these proposals are flawed because we cannot escape from our socio-historical tradition: we need a vocabulary for describing our experiences and we need guidance for assessing what kinds of processes reliably induce true belief. Given any suggested way of reconstructing our empirical knowledge, it's possible to imagine an alternative socio-historical tradition that would commend alternative ways of describing experience and would insist that the knowledge-generating processes be reliable against a different background of alternatives. The idea of showing the tradition-independence of empirical knowledge fails.

Of course, my main concern here isn't with empirical knowledge. Rather I want to emphasize the importance of the idea of the tradition-independence of a priori knowledge in the philosophy of mathematics. Consider the famous passages in which Frege introduces the notions of a priori and a posteriori as he understands them; these notions, he tells us, concern "the justification for making the judgment". In introducing the notions in this way, and in posing the problem of the foundations of mathematics as he does, Frege envisages a route to the propositions of mathematics (primarily arithmetic) that is tradition-independent, and those who debate the foundations of mathematics after him follow him in this. Suppose that the conception of a priori knowledge employed in these discussions were just the Weak conception. Then there are possible lives, given which processes that would normally warrant belief in various mathematical propositions would fail to do so. Now imagine a historical tradition whose members have such experiences in the generation that precedes ours. There are two possibilities: in socializing us they either respond to the subversive experiences by explicitly identifying certain processes as unreliable (in the way that we've learned to mark off certain kinds of perceptual situations

as unreliable) or they do not. If they do, then we are not warranted in believing parts of mathematics on the basis of the processes, any more than someone who has been told about mirages is warranted by his perceptions in believing that there is an oasis in the distance (even though there may be one there). If they do not, then we are still not warranted, for our epistemic situation is akin to that of people reared in a community of dedicated clairvoyants who ignore evidence that their chosen methods are unreliable. Either way, we who live in this hypothetical tradition are not warranted by processes that, as things stand (at least, so the foundationalist story goes) warrant us in our mathematical beliefs. Hence, the reconstruction of our knowledge turns out to be tradition-dependent.

On the other hand, if the Strong conception is adopted, then not only is this argument blocked but we see why Frege and his successors can suppose that they are offering a reconstruction of our mathematical knowledge that is independent of the social and historical circumstances of its genesis. They envisage displaying the structure of processes that would be available and which would warrant belief whatever our past experiences and whatever the past experiences of those who socialized us – any concerns about the reliability of these processes can be overridden by us, that is the point of (b). As I noted earlier, failure to insist on (b) results in the historical contingency of a priori knowledge.

Perhaps you may think that giving up the tradition-independence of a priori knowledge is not too bad. I claim, however, that this diverges from the classical idea of apriority in important ways

Let me offer you my preferred story about mathematical knowledge. I think that mathematical knowledge began with simple experiences of manipulating objects in the world. Over two millennia ago, some of our intellectual ancestors systematized the scraps of knowledge previously obtained by introducing postulates and definitions. Then, in the Renaissance, with the rise of early modern science, the community of inquirers instituted a division of labor. Some members of the community were to work on developing new mathematical concepts and principles, on the basis of what had already been achieved; others were to draw from the languages provided by the first group, formulating hypotheses about nature and devising experiments to test them. The institutionalization of that division has bequeathed to the language-makers of our day, who typically inhabit mathematics departments, a power to propose concepts and principles and to articulate those concepts, so long as they follow the rules of mathematical methodology, extending the mathematical traditions in which they have been schooled. They do not make observations or do experiments; they think and scribble. Their peers sometimes count what they produce as additions to knowledge, and such additions become available to their colleagues who do grub around in the world (or in the artificial world of the laboratory). Given the institutional backdrop, given the processes of socialization that have provided them with their starting points, their thoughts lead them to mathematical knowledge. I see no reason to doubt that the processes that warrant their beliefs often satisfy condition (a). But, of course, the warranting power is tradition-dependent.

The dependence on tradition connects contemporary mathematical knowledge with experience in two distinct ways. First, there is the obvious point

that the ultimate starting points lie in those scattered perceptions that began the whole show (and, as I've noted elsewhere, there may be other places where experience enters in the subsequent development of mathematics). Second, and less evident, the warranting power of the contemporary processes depends on the division of labor and the long sequence of experiences that have warranted our ancestors, and now us, in making that division. We have learned, from experience, that having a group of people who think and scribble, who proceed to extend and articulate mathematical languages in the ways that mathematicians find fruitful and who provide resources for empirical science is a good thing, that creating this role promotes our inquiry.

The tradition-dependence of contemporary mathematical knowledge indicates an important kind of dependence on experience. Philosophers beguiled by the obvious fact that mathematicians generate a lot of knowledge without any recourse to observation and experiment, and concerned to save the apriority of mathematical knowledge by weakening the conception so that it only demands (a), have failed to see that the warranting power of the processes of thought they take to underlie mathematical knowledge depends on the experiences of those who came before us in the mathematical tradition. Champions of the Weak conception save the apriority of mathematics by proposing, for example, that our current grasp of arithmetical concepts gives us a warrant for belief in the elementary truths of arithmetic, a warrant that is available given any sufficiently rich experience, but which, they concede, is defeasible. In emphasizing the tradition-dependence of such warrants, I've tried to transcend the inadequate individualistic epistemology that demands that we offer a synchronic reconstruction of human knowledge in which every item of knowledge is either categorized as empirical (grounded in the knower's own experiences) or a priori; if those were our options, then, I agree, classifying the knowledge as a priori would be preferable. But we ought to reject the individualistic framework, recognizing explicitly that our knowledge is dependent on the tradition in which we stand, and that our a priori knowledge, in the Weak sense, depends on the experiences of historical figures who have played important roles in past inquiry. Once we see this, we'll recognize that the issue isn't one of apriorism versus empiricism, but of apriorism versus historicism, and here the interesting question is whether one can find, for logic, mathematics, or anything else, some tradition-independent warrant, something that will meet the requirements that Descartes and Frege hoped to satisfy – in short, something that will answer to the Strong conception. Those who settle for the Weak conception have preserved the label 'a priori', but, to adapt a nice phrase of Michael Devitt's, they have offered us fig-leaf apriorism.

I don't offer this as a claim that the Strong conception captures what epistemologists have always understood about apriority. It seems to me that the discussions of the past decades have made clear how intricate and complex the classical notion of the a priori is, and that neither the Strong conception nor the Weak conception (nor anything else) can provide a coherent explication. In the end, it doesn't matter much whether we declare that mathematics isn't a priori (cleaving to the Strong conception) or whether we argue that mathemat-

ics is a priori (on the basis of the Weak conception). The important point is to understand the tradition-dependence of our mathematical knowledge and the complex mix of theoretical reasoning and empirical evidence that has figured in the historical process on which current mathematical knowledge is based.

VII

I want to extend these points by looking at two perceptive recent discussions of classical claims about the a priori. Both Hartry Field and Michael Friedman have offered defenses of the apriority of types of propositions traditionally favored with this status, defenses that seem to take into account many of the points I've been urging. Field is concerned with the apriority (or "aprioricity" as he calls it) of classical logic. Friedman wants to argue for the apriority of mathematics. In both instances, the defense takes the same form, in that beliefs in the pertinent propositions are supposed to be needed for people to perform certain kinds of cognitive tasks – engage in any assessment of evidence (Field), or formulate and assess the kinds of theories that we take to represent the pinnacle of our empirical knowledge (Friedman). I think this represents a confusion of the sort indicated in section I, between apriority and a quite different idea, the notion of a propositional precondition, and that the confusion makes the Weak conception seem more attractive than it actually is.

Field starts by formulating the issue in terms of evidential systems. Roughly, an epistemic system for a person is a system of rules that idealize that person's belief-forming and belief-retaining behavior: as I understand it, the rules are supposed to embody the person's own conception of evidence. A proposition p is weakly a priori for a person just in case the evidential system licenses belief even in absence of any particular sensory inputs and strongly a priori if, in addition, the rules don't license retraction of belief on the basis of any sensory inputs. The difference between these conceptions is parallel to the difference between the Weak conception and the Strong conception, although Field wants to rule out certain kinds of sensory inputs as not pertinent to questions about apriority. (These include the sorts of experiences I've invoked to discredit the apriority of mathematics, so it will be important, later on, to come to terms with his efforts at restriction.) So far the account of the a priori is highly subjectivist, since we can envisage people using all sorts of evidential systems and, having, in consequence, quite different sets of a priori propositions. Field is explicit about this, and proposes that the question is "whether or not it is a good thing to have an evidential system that licenses the beliefs in a (weakly or strongly) a priori fashion". Once he has posed the issue in this way, he can defend the apriority of logic in two steps, first by showing that our evidential system does license belief in logic without regard to sensory input, and second by arguing that our evidential system is a good thing.

If the aim is simply to establish the weak apriority of logic, then the first step seems trivial, for, if we understand evidential systems as Field does, then surely our evidential system does license endorsement of the principles of classical

logic in the absence of any particular sensory inputs. In pursuing a stronger conclusion, Field points out that we have no idea how to make sense of evidential relations involving the proposed rivals to classical logic, and that all the well-known reconstructions of measures of empirical evidence license our acceptance of classical logic, come what may. In my terminology, classical logic (or at least, its most central tenets) consists of propositional preconditions for our assessing empirical evidence in the way we do.

Turning now to the second step, Field doesn't make what seems to me a rather obvious point, namely that confidence in the progressiveness of inquiry ought to incline us to think that our evidential system is a good one. As we reflect on the history of science, it does appear that we've learned a lot about the world and we've learned a lot about how to learn about the world, and, while admitting our fallibility on all counts, we ought to endorse the view that we have a good evidential system. Field does, however, offer a related, forward-looking, point, noting that the possibility of a future system that would be superior to ours doesn't undermine the claim that our system is a good one. An important part of his case here is that evidential systems aren't correct or incorrect, so that we can't insist that the apriority of logic turns on whether our belief in classical logic is licensed by the right evidential system.

I'm more inclined to factualism in epistemology than Field is, but, for present purposes, I'm going to sidestep those thorny questions. The main point on which I want to insist is Field's divorce between strong apriority (in his sense) and the possibility of a priori knowledge (in the Strong sense). Assume that it's true that logic is weakly a priori and that logic is a set of propositional preconditions for the assessment of evidence. What does that show about our knowledge of the laws of classical logic?

Nothing. I believe the laws of classical logic, in part because I was taught them, and in part because I think I see how those laws are used in assessing evidence. But my belief could easily be undermined by experience. Perhaps I could be shown that the laws I've accepted really don't reconstruct the inferences that are made in the parts of our knowledge we esteem most (as Frege and Russell argued that classical Aristotelian logic was inadequate to reconstruct mathematical proofs); although my evidential system really does license the laws, it doesn't follow that I have to recognize this, or that I couldn't be brought to think that my system ought to be reconstructed by favoring different laws. In the history of inquiry human practice has often been at odds with the theoretical principles that people take to govern the reasoning in which they engage, and, even if that's not true of us, there are lots of ways in which our identification of logic might be disturbed.

Second, as Field acknowledges, we might modify our practice (although, as he points out, it's not easy to give substance to this idea). It's conceivable that future investigators may discover a new way to think about evidence and inference that transforms our approaches to these concepts even more radically than the modifications achieved in the wake of modern mathematical logic. There are two versions of this scenario, one in which our principles of logic turn out not to be correct (perhaps they involve faulty conceptualizations that

our successors transcend) and one in which we are seduced into believing that they are incorrect by the apparent successes of a rival community of investigators (these investigators convince us to adopt a different evidential system that is not actually as good as our own, despite appearances). Neither way of articulating the scenario is at odds with the claims that our evidential system licenses the laws of classical logic, irrespective of sensory input, and that our evidential system is a good one. Weak apriority of logic is sustained.

But it's easy to see how either scenario would affect the view that we know the laws of logic a priori in the sense of the Strong conception. Imagine a community that lives through one of the envisaged transitions. After the presentation of the powerful, or apparently powerful, rival system, the community splits. Most become converted to the new way of doing things, and they no longer believe the laws of our logic, but a small group of traditionalists persist in the old ways. Notice that they believe the laws, they endorse an evidential system that licenses the belief, and their system is a good one. Nonetheless, their beliefs are based on ignoring or misevaluating the apparent advantages of the revisionist way of doing things. Under these conditions, we can't see them as anything other than blind dogmatists, and, in consequence, we shouldn't view their beliefs as warranted. So, even though classical logic may contain (even comprise) propositional preconditions of experience (or evidence), it's not the case that classical logic can be known Strongly a priori.

Now it may quite reasonably be suspected that my argument against Field depends crucially on ignoring a distinction he makes quite early on in his discussion, the distinction between primary and secondary undermining evidence. The latter "does not primarily go against the claim being undermined but against the claim that we knew it a priori", and Field proposes that strong apriority doesn't require the ability to survive secondary undermining evidence. Field might contend that the only debunking scenarios involve secondary evidence.

When we probe this contention, we see that the distinction is much less clear than it initially appears. Consider a standard undermining story. Our descendants come across a group of inquirers who seem, by all our lights, extraordinarily successful: their predictions even of complex phenomena are quite staggering. They try to teach us their evidential system and we find it very hard to understand (perhaps in much the way some students find logic alien and tough going), but some of us are able to follow their recommendations about inference and those people obtain apparently wonderful results. Maybe they offer us "proofs" of difficulties within our own system. These might be more complex versions of the "proofs" that $1 = 2$ that sometimes baffle schoolchildren, and, I'll assume, diagnosing the trouble is beyond our abilities. Does this hypothetical sequence of experiences furnish us with primary or secondary undermining evidence? If it's primary, then Field's distinction is unavailing with this sort of case. If, on the other hand, it's secondary, then one has to restrict the notion of primary evidence quite narrowly. After all, we might say that there's indirect primary evidence against the truth of our logical principles and that this is furnished by the demonstration of success of a group of people who proceed in ways that support the idea that they are using different principles, as well

as by the (presumably fallacious) "proofs". If one denies that there can ever be primary evidence where the testimony of others is involved, then we have extraordinarily little primary evidence for anything – that's the lesson of the socio-historical view of knowledge. I think that Field's distinction is influenced, at bottom, by the notion that there is a sort of evidence that is available to individuals independently of their training and of the history of the society in which they live. Once we understand how little (if anything) would qualify for that role, then I don't think one can exclude the debunking scenarios without robbing the notion of primary evidence of any interest.

Although Field is much less friendly to apriorism than Friedman, there's a striking convergence in their recent strategies. In a penetrating recent discussion of philosophical naturalism, Friedman follows Reichenbach in distinguishing two notions of the a priori: "necessary and unrevisable, fixed for all time, on the one hand, "constitutive of the object of [scientific] knowledge", on the other." The real hero of the story is Carnap, who "brought this new, relativized and dynamical conception of the a priori to its most precise expression". According to Carnap, our formulations of hypotheses and our assessment of them presuppose a mathematico-logico-linguistic framework, and, within any such framework, the mathematical principles play a constitutive role, so that they cannot be questioned while we remain with the framework. Of course, as Carnap took pains to point out, we can change frameworks, inscribing a different set of principles in our practice of formulating and testing empirical hypotheses. In changes of framework, however, we appeal not only to empirical evidence but also to specifically mathematical and even philosophical considerations. Instead of the holistic vision of evidence familiar from Quine, Friedman offers an alternative

... the picture of a dynamical system of beliefs, concepts, and principles that can be analyzed, for present purposes, into three main components: an evolving system of empirical natural scientific concepts and principles, an evolving system of mathematical concepts and principles which frame those of empirical science and make their rigorous formulation and precise experimental testing possible, and an evolving system of philosophical concepts and principles which serve, especially in periods of conceptual revolution, as a source of suggestions and guidance in choosing one scientific framework rather than another. All of these systems are in continual dynamical evolution, and it is indeed the case that no concept or principle is forever immune to revision. Yet we can nonetheless clearly distinguish the radically different functions, levels, and roles of the differing component systems. In particular, although the three component systems are certainly in perpetual interaction, they nonetheless evolve according to their own characteristic dynamics.

There is much here with which I agree. What puzzles me most is why this should be thought of as any kind of rehabilitation of apriority.

For present purposes I'm going to ignore one of Friedman's component systems (the body of philosophical lore, whose role in his argument is much less clearly delineated) and consider the image of the interplay between mathematics and natural science that he offers us. Mathematicians supply some of the ingredients out of which frameworks are constructed, and scientists draw on their

offerings to formulate successors to previous empirical theories in the light of observations and experiments. Now there are two quite distinct ways to think about the way in which the mathematical tools are fashioned, or, in Friedman's idiom about the dynamics according to which this component system evolves.

According to one of these, mathematical knowledge grows by the kinds of processes I've discussed in the later chapters of *The Nature of Mathematical Knowledge*: mathematicians try to devise ways of responding to unanswered questions, they develop generalizations of methods and concepts that have previously been introduced, they attempt to find rigorous replacements for forms of reasoning that appear to yield correct conclusions but do not accord with prevailing standards – and, in doing all this, they propose definitions and axioms that systematize previous mathematics. The ultimate roots of their practice lie in experience, and, from time to time, connections with the empirical sciences are reformed, and particular parts of mathematics are inspired by the needs of particular lines of investigation into natural (or social) phenomena. Hence, there is genuine co-evolution between mathematics and science. Yet, above all, the entire strategy of licensing a group of people to develop languages and principles in this way – essentially by thinking about what has been achieved in the past and how to generalize and systematize it further – rests on our appreciation of its fruitfulness for inquiry. The knowledge of contemporary mathematicians may be proximally produced by their reflections on what they have absorbed from the past, reflections that do not depend on any specific sensory input, but it is ultimately dependent on the collective experiences of the tradition in which they stand.

If this is Friedman's picture, then he and I are in fundamental accord. We can even agree to call mathematical knowledge "relatively a priori", or, as I would prefer, "proximally a priori", so long as we are clear that this Weak conception diverges from traditional ideas about a priori mathematical knowledge. For the tradition sees the dynamic of the mathematical subsystem in very different terms. Mathematicians are supposed to have ways of warranting their beliefs that stand outside the historical process, that are independent of tradition. It would be possible to try to articulate Friedman's vision by adopting this idea, insisting that mathematical knowledge is a priori in the sense of the Strong conception, although I think that that would be at odds with his point about the Carnap-Reichenbach insight. Any such program would return mathematical epistemology to the venture of seeking processes that can serve as tradition-independent warrants – a venture I believe to be fruitless. Field's efforts to secure the apriority of logic seem to me to be an ingenious last-ditch attempt to bind the relativized notion of apriority (the Carnap-Reichenbach, "constitutive", conception) to the more traditional notion, by arguing that the framework-constituting principles are immune to empirical evidence. I've argued that those efforts fail. The proximal a priori is the best one can hope for.

Let me set this in the context of my main point. Traditionally apriorism has been opposed to radical empiricism, and, in my earlier discussions of the issues, I tended to endorse that opposition. But it now seems to me that apriorism and empiricism are both best understood as rival versions of attempts to

give a tradition-independent conception of human knowledge. The most fundamental feature of my attack on apriority is my attempt to debunk tradition-independence. If I interpret them correctly, Field thinks he can deploy the notion of a propositional precondition of experience to defend the tradition-independence of our knowledge of logic, while Friedman thinks that this notion will retain the significance of classical theses about apriority, even when we admit the "relative" (tradition-dependent) character of apriority. In my judgment, neither of these last-ditch efforts succeeds. I recommend declaring a truce on the apriorism/empiricism debate on the grounds that logic, mathematics and (as Section V suggested) whole chunks of other disciplines count as a priori in the Weak sense. That truce should be coupled with a clear understanding of the places in which the Weak conception departs from classical lore about apriority, agreement on the tradition-dependence of our logical and mathematical knowledge, and a resolve to explore the complex ways in which experience has figured in the genesis of our current logical and mathematical knowledge.