

Computer Proof, Apriori Knowledge, and Other Minds: The Sixth Philosophical Perspectives Lecture

Author(s): Tyler Burge

Source: *Philosophical Perspectives*, Vol. 12, Language, Mind, and Ontology (1998), pp. 1-37

Published by: Ridgeview Publishing Company

Stable URL: <http://www.jstor.org/stable/2676139>

Accessed: 11-04-2017 02:09 UTC

REFERENCES

Linked references are available on JSTOR for this article:

http://www.jstor.org/stable/2676139?seq=1&cid=pdf-reference#references_tab_contents

You may need to log in to JSTOR to access the linked references.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at

<http://about.jstor.org/terms>



Ridgeview Publishing Company is collaborating with JSTOR to digitize, preserve and extend access to *Philosophical Perspectives*

COMPUTER PROOF, APRIORI KNOWLEDGE, AND OTHER MINDS*

The Sixth *Philosophical Perspectives* Lecture

Tyler Burge
University of California, Los Angeles

Almost two decades ago there was a flurry about the significance of the computer-proved Four-Color Theorem. Humans had programmed the computer, of course. But it had run a proof that was too long for any human to carry out or check. The full proof could be checked only by other computers. Nevertheless, nearly all mathematicians conceded that the theorem had been proved.¹ The status of the theorem has not changed substantially since then. I will be interested in the situation only as thus described. I want to discuss a philosophical question that arises from the supposition that we are ineliminably dependent on computers for some of our mathematical knowledge.

The question concerns the justification for a person's believing, indeed knowing, a theorem on the basis of a computer proof. I want to ask whether reliance on a computer must in itself prevent one's warrant for believing a theorem from being apriori.

The dominant philosophical view is that reliance on a computer in itself makes a person's justification empirical—prevents it from being apriori. Discussing an analogous case (without reference to whether the computer's work could be replicated by humans), Saul Kripke argues that such beliefs are a posteriori:

No one has calculated or proved that the number is prime; but the machine has given the answer: this number is prime. We, then, if we believe that the number is prime, believe it on the basis of our knowledge of the laws of physics, the construction of the machine, and so on.²

Thomas Tymoczko writes:

This appeal to computer, whether we count it strictly a part of a proof or as a part of some explicitly non-proof-theoretic component of mathematical knowledge, is ulti-

mately a report on a successful experiment. It helps establish 4CT...on grounds that are in part empirical.³

These views are plausible. But the issue is more complicated than they indicate. I think that their suggestion that knowledge of these results *must* rest partly on an empirical justification is problematic and probably mistaken. I want to explore the situation in a Socratic spirit, aiming less to settle the matter than to identify and broach some epistemological questions.

The case is worth discussing for its own sake. Issues regarding unsurveyable proofs and issues regarding the reliance on computers or on interpersonal communication in mathematics are interesting in themselves. But a co-equal motive for discussing the case is that it provides a context for exploring a distinction that is fundamental to understanding rationalism as an epistemological position.

For most of this century, various forms of deflationary empiricism have dominated thinking about knowledge. In my view, this dominance is poorly grounded and will collapse. For all empiricism's success in criticizing a procession of unfortunate metaphysical views and in championing the methods of natural science, it fails to yield a credible, complete account of our knowledge of logic, mathematics, self-knowledge, ethics, and several other parts of philosophy. The dominance of empiricism has led to complacency regarding empiricism's vulnerabilities and ignorance regarding rationalism's resources.

It is still too little recognized that rationalism is not about unrevisability, infallibility, indubitability, or innateness. It is about the force of epistemic warrant. It claims that in some cases the force of our warrant for believing (or knowing) certain propositions or for relying on certain practices or capacities derives from intellectual understanding or reason, not from sense experience. The phrase "force of our warrant" should be highlighted. Rationalists need not (and many classical rationalists did not) claim that any of our beliefs are independent of sense experience *tout court*. Perhaps all of our thinking depends on having had sense experience. Rationalism claims rather that sense experience does not contribute to the normative or justificational force carried by some warrants. So arguing that a belief "depends on" sense experience does nothing in itself to support an empiricist epistemology.

The issue centers on distinguishing epistemic or justificational dependence from various other sorts of dependence—causal or psychological dependence, or dependence for grasping intentional content or for learning. If one does not recognize the *prima facie* possibility of drawing this distinction, one cannot even begin to reflect on the issues between rationalism and empiricism, or to see what rationalism is driving at. The point has been emphasized by nearly every major rationalist.⁴ But it is repeatedly forgotten in current discussions. I shall explore the distinction between dependence on the senses for acquiring putative information and dependence on the senses for warranting our understanding and acceptance of the information thus acquired. I will be developing a rationalist line. But a more fundamental aim is to make progress on understanding the issues

surrounding computer proof, with special reference to this distinction. Understanding the distinction is more important than the position on the special case that forms the topic of this paper.

I

I begin by explaining what I mean by “apriori”. I take apriority to be a feature primarily of justification or entitlement. A warrant (either a justification or an entitlement), is *apriori* if neither sense experiences nor sense-perceptual beliefs are referred to or relied upon to contribute to the justificational force particular to that warrant.⁵ A person’s knowledge is apriori if the knowledge is supported by an apriori warrant that needs no further warrant for the knowledge to be knowledge.

Let me highlight some features of this conception of apriority.⁶ The distinction between justification and entitlement concerns two types of epistemic warrant. Both justifications and entitlements are epistemic warrants with rational justificational force. But entitlements need not be understood by or even accessible to the individual subject, whereas justifications, in my narrow sense, involve reasons that individuals have and have access to. Entitlement is my partially externalist analog of the internalist notion, justification. I will leave open exactly what may count as an entitlement, although I do not believe that *any* entitlements are *mere* matters of reliability. What is important is that with an entitlement, a *full* reason, or a full expression of the warrant associated with a state or capacity need not be available to the individual.

We are entitled to rely on our capacities for perception, memory, interlocution, deductive and inductive reasoning. Children are entitled to rely on particular perceptual beliefs even though they could not understand why. Perhaps only philosophers can explain why. Justifications, in my narrow sense, are available on reflection to the justified individuals. Justifications may be self-sufficient premises, or reasons that a person could “in principle” call up.

That a warrant for a belief is apriori does not entail that the belief is self-evident, infallible, indubitable, innate, unrevisable, or even unrevisable by *empirical* considerations. Apriority concerns the nature of the rational support for an attitude, not the nature of its vulnerability to criticism.

A non-demonstrative justification can be apriori. Arguably, much non-demonstrative reasoning in mathematics is apriori: neither sense experiences nor sense-perceptual beliefs need contribute to the justificational force involved in such reasoning. The reasoning explains or proceeds inductively from judgments that are fundamentally mathematical or intellectual, not sense-perceptual, in their justificational underpinnings.

Sense experience may be psychologically necessary for an acquisition or warrant of a belief, without contributing to the force of the warrant—hence without making it empirical. The role of sense experience in the psychology and acquisition of belief must be distinguished from its role in contributing to the normative force associated with the belief’s warrant.

In deciding whether a belief is apriori warranted, it does not suffice to determine whether it is empirically warranted. A belief can have empirical as well as apriori justification or entitlement. The issue is always whether there is justification or entitlement that has justificational force independently of any force contributed by reliance on sense experiences or sense-perceptual beliefs.

I shall let these remarks suffice for now to provide a working understanding of my conception of apriority. But my stage-setting is not complete. I shall make four large, controversial assumptions. I stand by these assumptions. But I make them not primarily to support a conclusion but to isolate the issues I want to discuss. I hope that the discussion will be of some interest even to those who do not share the assumptions. Here I make them for the sake of argument.

First, I assume that individuals' knowledge of pure mathematics, resting on specifically mathematical understanding or reasoning, is ordinarily apriori.⁷ Although this assumption is larger than the thesis under discussion, its role here is just to help us focus on whether knowledge of computer-proved theorems *must in itself* add an empirical element to the warrant for mathematical knowledge.

Second, I assume that the computers that are used in carrying out the relevant proofs do not have minds and are not autonomous thinkers. They can do too little to relate their operations to reality to count as having autonomous meanings or thought contents. For example, they have no perceptual abilities and do nothing for themselves. Their manipulation of mathematical symbols is meaningful mathematical activity only insofar as human beings have programmed the machines to manipulate contentful mathematical symbols. Whether robots could have autonomous thought is a question I leave open.

Third, I assume that we, including mathematicians among us, can obtain mathematical knowledge from others, even when we do not know a proof. There are positivistic views, still voiced sometimes, according to which one cannot know (or even, on some views, understand) a mathematical proposition unless one knows a proof of it. This seems to me quite unbelievable as applied to knowledge in an ordinary unstressed sense. Many people know the Pythagorean theorem who have never seen a proof of it. Mathematicians know theorems from other mathematicians, even though they have only the vaguest idea of how the proofs go. These are cases where someone among the individual's sources knows a proof. But even this does not seem necessary. Newton knew simple truths of the calculus, even though they were not self-evident and even though no one, including Newton, knew a proof at the time. The positivist, and intuitionist, insistence on proof for mathematical knowledge is simply a mistake.

Fourth, I assume, what I have argued elsewhere, that although commonly empirical, knowledge that an individual obtains by being told a proposition by another person, where the individual's warrant resides in the interlocution, *can be apriori*.⁸ If one's source knows a proposition apriori and there is no reason to doubt the source, one can in certain cases know it apriori by accepting the source's word.

This is the least familiar of my assumptions. Again, I make it to isolate distinctive aspects of the epistemology of reliance on computer proof that do not arise in ordinary interlocution. Because of its unfamiliarity, I want to elaborate it at least briefly.

The basic idea is that rational acceptance of an interlocutor's word is an apriori default position. The seeming intelligibility of another's talk gives one apriori prima facie warrant to presume that the other is a resource for rationality. Resources for rationality are, other things equal, to be believed. The presumption of another's credibility can be defeated by special considerations, including empirical considerations. But in the absence of reasonable doubt it is a rational starting point—independent of empirical support. Sometimes, acting in accord with this default entitlement, together with the fact that one's source is knowledgeable regarding the information being passed on, suffices to provide one with apriori knowledge. In such cases, the recipient's apriori prima facie warrant can reside fundamentally in the interlocution, not in an independent and sufficient warrant that the recipient has for the communicated knowledge.

Of course, most knowledge preserved through interlocution is empirical. When a recipient is told something, there are a multitude of possible empirical sources of warrant. The original knowledge passed along might itself be empirical. The recipient may need empirical grounds to supplement his apriori default warrant—because there may be reason to doubt the interlocutor's competence or veracity. The recipient may need empirical observation or empirical reasoning to justify an understanding of the interlocutor's utterance. The recipient may have reason to believe that the topic (like politics or philosophy or many specialized subjects) is one on which one cannot accept just anyone's word; and empirical reasons may be needed to bolster the credentials of the interlocutor or the plausibility of what he says. It is only in special cases that apriori knowledge can be preserved through interlocution. What I am assuming, with some background of argument, is the mere possibility of such preservation. The explication of apriori knowledge that I gave requires that the knowledge be supported by an apriori justification or entitlement that needs no further justification or entitlement for it to be knowledge. Knowledge gained by interlocution raises special problems. When one depends on an interlocutor for knowledge, one's knowledge depends partly on one's own entitlement for accepting the interlocutor's word and partly on the knowledgeability of the interlocutor.⁹ More particularly, the recipient depends on there being in the chain of sources sufficient justification or entitlement to underwrite knowledge—or else indication of the existence of a method of justification sufficient to underwrite knowledge. For the recipient's entitlement to rely on interlocution is never by itself sufficient to underwrite the recipient's knowledge.¹⁰

So in interlocution we distinguish two bodies of epistemic warrant: (i) the recipient's *proprietary warrant* for a belief—that is, the reasons available to him together with his epistemic entitlements for holding the belief; and (ii) the *ex-*

tended body of warrant for a belief—which includes not only the recipient’s proprietary warrant, but those warrants for the belief that are possessed or indicated by interlocutors on whom the recipient depends for his knowledge (though not for his proprietary warrant). The recipient “depends” on interlocutors’ justificational resources for his knowledge in the sense that if the interlocutors had lacked or failed to indicate the existence of these justificational resources, the recipient’s warranted true belief would not have been knowledge. A person’s knowledge based on interlocation is apriori only if some sufficient part of the extended body of warrant (including also an appropriate part of the subject’s proprietary warrant) for the knowledge is apriori. There must be a body of warrant or justificational resources within the extended body of warrant that suffices for knowledge and that is itself apriori.

II

I turn now to accounts of why reliance on computers for proof of a theorem must be empirically justified. Tymoczko writes,

Some people might be tempted to accept appeal to computers on the ground that it involves a harmless extension of human powers. On their view, the computer merely traces out the steps of a complicated formal proof that is really out there. In fact, our only evidence for the existence of that formal proof presupposes the reliability of computers.¹¹

Tymoczko adds that our knowledge that computers are reliable can only be justified empirically.

One difficulty with this argument is that the key word “presupposes” is unclear. The question is in what form an assumption about the reliability of the computers is an indispensable part of the warrant for believing in the existence of a proof. In a certain sense our ordinary formal arguments for mathematical results “presuppose” the reliability of our brain processes. If our brains were relevantly defective, a faulty proof could seem to be sound. We can know that our brains are reliable only by empirical means. This hardly shows that all our formal arguments in mathematics are empirical. Although we depend on the reliability of our brains, the assumption of their reliability need not be a part of the warrant for our mathematical result. The question is how our reliance on a computer’s reliability is different from our reliance on our brains’ reliability. I believe that the reliance *is* different. But do the differences force the warrant for relying on a computer to be empirical?

A similar question arises about Kripke’s claim that one believes a computer’s saying that a number is prime only “on the basis of” knowledge of the laws of physics and the construction of the machine. It is true that we could have gotten the wrong answer if the physical world had been less regular or if the machine had been physically defective. And it is true that we may and do support our belief

with empirical information about the machine. But it does not follow that a justification of our belief must rely for its justificational force on these empirically known physical facts. To be justified in believing the machine's mathematical results, do we *have to* include reference to the physical mechanism underlying the proof, when we do not need to do so in giving ordinary proofs? If so, why?

An obvious answer is that when we have a proof, the proof suffices for our justification. Unless there is some actual defect in the proof, there is no defect in the justification. So our knowledge that our reasoning depends on our brains is not needed in carrying out an apriori mathematical justification. But in the situation we are considering, we do not have a proof—at least not one of our own that we can check fully. We rely on the computer's having carried out the details. So the justification might appear to have to make reference to the computer's physical reliability. Perhaps this is the relevant difference between the two cases.

The difference is relevant. But although we do not have a proof "of our own", this does not suffice to explain why our belief in the computer-produced theorem must be empirically justified. For in accepting information from other people, we often lack our own justification for the acquired beliefs. We presume that the source has, or indicates the existence of, justificational resources; we need not have those resources ourselves. Yet we do not always need to know that our source is reliable in order to be warranted in accepting what the source says. It is enough if we have no reason to think the source unreliable. Acceptance is a default position, a rational starting point that needs no empirical support in the absence of reasonable doubt. One's fundamental entitlement to accept what one is told is, in the absence of reasonable doubt, apriori. This was our fourth assumption above.

Thus on our assumptions, one can have apriori knowledge of a proposition one has received from another source, even though one cannot oneself provide a justification sufficient in itself to make one's belief knowledge. One relies on there being justificational resources in one's chain of sources. The extended body of warrant that enables one to have knowledge might be apriori. People who accept the Pythagorean theorem (or comparable theorems in arithmetic) on the basis of a diagram and the word of another, but who themselves cannot produce a proof of it, are often in this position.

The case I cited in which one could have such knowledge is one in which one lacks reasonable doubt about what one's source says. So our questions have become: Why must there be reasonable doubt about the computer, so that its outputs cannot fall under one's general right to accept the word of others? And why must such doubt force empirical considerations to underwrite belief in the computer's offering? I will consider several answers to these questions.

To be warranted in believing the computer's result, we do need reasons to supplement our entitlement to believe what others tell us. The theorem's difficulty and the fact that no one has checked the proof force some supplement. But I want to postpone discussing this issue, to focus first on this narrower question: Is the unsurveyability of the computer's deduction sufficient ground in itself for

taking belief in the computer's result to be empirical?¹² The answer is "no".¹³ An unsurveyable deductive argument is no *more* inherently empirical than a non-demonstrative argument in mathematics.

Much justified mathematical belief is underwritten by non-demonstrative reasoning: Newton's belief in, indeed knowledge of, the elementary truths of the calculus, before they were given mathematical explications, is an example. Our belief in the consistency of arithmetic seems thoroughly warranted; in fact I think it constitutes knowledge. But no proof of it adds significantly to the ground for our belief. Zermelo's belief in the axiom of choice, new axioms in descriptive set theory, and Church's thesis also count as examples of warranted belief, though the latter two cases may not amount to knowledge. It was reasonable to believe these propositions in the absence of proof, and in the absence of self-evidence. The relevant considerations were mathematical. There appears to be no indispensable reliance for justification on sense experience, or on the role of the arguments in physical theory.¹⁴ The form or length of an argument does not determine whether its justification must depend, in the relevant way, on sense experience.

Let us try a new tack. What role does the source's being a computer, rather than a person, play in our evaluation of the argument? Is this a source of reasonable doubt about the computer's output? Is it a necessary source of empirical elements in our justification? These issues are complicated. In accepting the word of others in normal discourse, we presume that the seeming-intelligibility of others' talk is an *apriori prima face* sign of the others' rationality. When we accept the word of another, we treat the other as a rational source. Lacking reasons to the contrary, we are *apriori prima facie* entitled to accept what others tell us. But by our third assumption, the relevant computers are not people. They do not have minds. They do not have warrants, strictly, because they lack beliefs. They are not autonomously rational.

To investigate whether our source's being a computer is a ground for reasonable doubt about its output—and a source of empirical elements in our justification—I will consider what about the computer can be seen as possible sources of error.

The programming language and the programming of the computer constitute possible sources of error—by far the most likely sources. But these are themselves versions of a mathematical language and a strategy for mathematical argument. The problems they introduce are not in principle different from problems inherent in any mathematical language or mathematical form of reasoning. Since we are assuming that ordinary mathematical reasoning is *apriori*, I shall ignore these aspects of the problem of relying on a computer to give us mathematical information. The language and program are expressions of the mathematical rationality of the programmers.

Of course, there are physical transitions in putting the program into the computer and in running the program. These are possible sources of error. The source might lie in some defect in the machine, or in some error of programming execution. Let us consider errors of programming execution. Errors in entering the

program into the machine seem analogous to errors of expression in ordinary mathematical argumentation. Proofs can go wrong because the author writes down a different symbol from the one intended. Although such slips are sources of error that can be checked for only by empirical means (one has to look at the written symbols to verify that they are the ones intended), this does not show that mathematical argumentation that uses written symbols, diagrams, and so on, is empirical. Since we are making the large assumption that ordinary mathematical reasoning is not empirical, I lay this matter aside here.¹⁵

Because of some physical defect, the computer may fail to carry out correctly the mathematical program that has been entered into it. As noted, an analogous point applies to human beings. Defects in others' brains may distort their sense for correctness. But we may presume that they are rational sources, or resources for rationality, absent reasonable doubt, because we may consider the content of their output without knowing or needing to know anything about the physical mechanisms that produce the output. The same point applies to interaction with computers. We take up their language, their mathematical offerings, and need not consider their physique.

There is this difference: Machines must be made. The designers must make them carry out rational procedures (make them computers) before anything can be presumed about their reliability as interlocutors. The computers we are discussing are not autonomously rational.¹⁶ To make a computer, the maker had first to view it as an object of empirical engineering. The maker could not simply presume that the machine carried out instructions according to the rational principles that the maker intended it to realize. The maker had to check and revise, through empirical experimentation, to make the machine into a computer capable of carrying out its intended functions. In the designing process, it does seem appropriate to see the machine's output as the result of an empirical experiment. But so far we have seen no reason to think that we must always join the designers in requiring empirical assurance that the machine is a reliable physical realizer of the relevant design intentions, as long as the machine is producing seemingly intelligible and credible information.

Empirical knowledge is available about the physical reliability of the computer. Having this knowledge may increase our confidence in the computer's offerings. But given that we have rationally intelligible and credible mathematical presentations from the machine, we can presume that we are being confronted with a resource for rationality, in the absence of specific reasons for doubt. Following a large part of the machine's proof can indicate the existence, form, and nature of a justification (proof) for a theorem. This remains true even though the machine does not itself "have" a justification.

The meaning of the machine's activities, and indeed the nature of its rational powers, are derivative from intentionality and rationality of rational beings. Such activities provide an amplification of the designers' and programmers' rationality. In understanding the machine's offerings, we have an *a priori* entitlement to rely on the rationality of these offerings, and the justifications backing these

offerings. If the justification indicated by the machine's output contains no premises or rules whose only justification is empirical, and supposing that we are apriori entitled to accept the machine's propositional offerings and proof-fragments as backed by rational procedures and rational sources, our extended body of warrant can be apriori. Again I invoke my assumptions about interlocution and about the apriority of the relevant mathematics.

III

The central epistemic issue is not that we are dealing with computers, but that the reasoning that we depend on is not only not ours but is known to be difficult. To be warranted in relying on the computer here, we need more assurance regarding its capabilities than can be gotten from knowing that it is a source of plausible mathematical presentations. Could this assurance be gained non-empirically? Let me approach this question through a parable.

Imagine that a mathematician (M) had computing power and speed equal to that of the relevant computer. Suppose that M had as much access to the details of the unconscious computations going on in her as the access that the programmers have to the computer's computations in the actual situation. Such computations are too long, fast, and complex to be brought fully to M's conscious mind and checked. But parts of the computations can be brought to mind. Suppose that M understands the principles of the unconscious computations that she carries out exactly as well as the programmers understand the nature of the program that they have entered into the machine. So M creates a proof strategy that is the same as the one the programmers create.¹⁷ M knows that she can carry out correct proofs unconsciously, since she has often checked the details of shorter, but still difficult proofs that she first carried out unconsciously. M's ability to reach sound conclusions unconsciously from her premises by way of consciously understood proof strategies is well-established. M's unconscious mathematical powers and track record exactly parallel those of the relevant computers.

To keep the parallels going, suppose that other mathematicians have computational powers similar to M's. Such mathematicians check each others' proofs by running fast unconscious calculations—in the way one computer checks another. Sometimes these checks can be consciously replicated by those mathematicians and by lesser humans, and they are found to be very reliable, though not infallible.¹⁸

Suppose that M tries a proof that is not consciously checkable. Any portion can be consciously checked. The strategy can be consciously checked. But the proof is too long to be fully gone over. M reaches a result and reports it. After the proof has been checked by M and similarly endowed mathematicians, who also cannot check all details consciously, the result is found to hold. Clearly, M and her colleagues would have reason to conclude that the theorem had been proved—

much as actual mathematicians have reason to conclude that computers have produced proofs that cannot be checked by humans.

I think that M's justification for believing her conclusion remains apriori. The full proof demonstration is not available to her, as a conscious, deliberative being. Her conscious justification for believing that the result has been proved contains a non-demonstrative element. She relies not only on her understanding of the premises and proof strategy, but on her unconscious mathematical powers. But, as I have noted, non-demonstrative elements in mathematical reasoning do not themselves render the reasoning empirical. The mathematical powers that M relies upon are an authoritative resource for rationality, analogous to other reasoning that M relies upon.

What is M's warrant for relying on her partly inaccessible, unconscious mathematical powers? I think that she has two warrants. One is an entitlement; one is an inductive justification.

She is apriori *entitled* to rely on those powers even lacking an articulateable justification. *She is entitled because those powers are in fact rational and she knows how to use them with sufficient understanding.* This is, I think, the form of our entitlement for all reliance on our rational faculties or our other resources for reason.

M's second type of warrant is an inductive *justification* for relying upon the relevant rational faculty. The induction is roughly, "I unconsciously proved t_1 , t_2 , ... t_m ; so I proved t_n ." I think that the knowledge of the induction base and the relevant type of inductive reasoning are intellectual, non-sense-perceptual, and apriori.

How is the induction base warranted? There are the provings of t_1 ... t_m . These are apriori warranted on our first assumption. Then there is the knowledge that she has proved them. This knowledge rests first on self-knowledge: a knowledge that one is thinking through the proof, together with knowledge that it is a proof. This self-knowledge is, I think, intellectual and does not rest for its justificational force on sense experience. Then, second, there is M's memory of her thinkings-through of these proofs. One function of such memory is the same as that of preservative memory, which makes possible thinking through any argument over time. The memory at issue here preserves not only the content of steps in arguments, but the past events of having reasoned through those steps and having recognized them as a proof. Such memory is fundamentally intellectual. Like the relevant self-knowledge, it does not rest on sense experience for its justificational force. It is a preservation of intellectual events and their content.

The memory of one's thinkings-through of past proofs also functions to support a continuing use of the first-person. I think that one has an intellectual, apriori warrant through direct memory to attribute past intellectual acts to oneself via the first-person concept. Of course, such warrant is defeasible, even empirically defeasible. But if such memory of one's past intellectual acts—which in the normal case seems not to involve any reliance on sense experience—, were not

warranted, one could not hold a rational deliberative point of view together over time.¹⁹

So the warrant for the combination of the mathematical knowledge of the proof together with the self-knowledge of one's proving the theorem and the memory preserving that self-knowledge, remains apriori. The preservative memory and relevant sorts of self-knowledge are underwritten by apriori, non-sense-perceptual entitlements; and the original mathematical knowledge is by hypothesis apriori.²⁰

But it is also plausible that the transition rules or norms of passage in the inductive inference that enables M to infer from her induction base regarding past proofs to the conclusion that she has proved the new theorem are also apriori. Such induction is as constitutive of being rational, to rationally processing information (including empirical information), as deductive inference is. So I believe that the induction as a whole is warranted apriori.

Defending all this in detail is beyond my scope here. I present it as plausible. It is certainly clear that in knowing that a proof has been carried out, M need not rely on knowledge of the physical laws governing her brain, or any measurements of the reliability of her *physical* capacities. I think that, at the very least, the empiricist has a tough road in showing that empirical sources of warrant *must* play a role in underwriting M's knowledge of the induction.

IV

Of course, computers are outside us in a way that M's powers are not outside her. We cannot access their results from a first-person point of view, as M could access her results. This remark raises three issues. First, we have to perceive physical events to know the results the computers express. Second, M has learned to use her mathematical powers, and has them in her, whereas we take over the computer's reasoning second-hand, regarding it from the third-person point of view. Third, to know that a computer can deal with difficult mathematical problems, to have some sense of the *degree* of its mathematical power, we must have used our sense capacities. Do these points show that our justification for believing that a computer has carried out an unsurveyable difficult proof is inevitably empirical?

All these issues have analogs in the epistemology of communicating with other people. I shall deal with the first two quickly, so as to concentrate on the third. As to our reliance on perception in knowing what a computer's results are: We do have to hear the words of others, just as we have to read the print-outs of computers. But perceiving the words need not be, and normally is not, part of the fundamental entitlement for understanding what intentional content they convey, or for understanding them as presentations-as-true. The words are like diagrams or symbols that help one see the point of a mathematical claim: They call to the recipient's mind the thoughts of the interlocutor that they express. There is no

other way. But perception of them does not normally enter into the recipient's entitlement for relying on understanding of them, or their assertive force. These issues are complex, but they are covered by my assumption that we can have apriori knowledge through interlocution.²¹

Second, although M has mastered the forms of her unconscious reasoning, that reasoning is no more available to her as a justification than the computer's transformations are available to us. We have to access through print-outs the computer's results and the individual passages of the proofs that we check, whereas M knows her results and passages of her reasoning "directly". But this reduces to the previous point about the role of perception in interlocution.

I turn now to the third issue, which I want to dwell on for the remainder of the paper. To have some sense of whether the computer has mathematical power sufficient to solve a difficult mathematical problem, we seem to have to rely on perception in a different way than that involved in simply understanding the print-outs of results. We have apriori prima facie entitlement to accept intelligible presentations-as-true, expressed by the print-outs. But the entitlement can be overridden if there is reasonable ground for doubting what one's interlocutor presents. The known difficulty of the relevant mathematical theorems seems to be prima facie ground for doubt or suspension of belief, in the absence of supplemental knowledge about the power and reliability of the interlocutor.

Suppose that I know that some proposition (say, Fermat's Last Theorem or the Four Color Theorem) has been alleged to be proved many times before, always mistakenly. If a stranger tells me that he or she has proved it, I have reasonable ground for withholding belief. As things stand, I am not justified in believing what I am told. The initial generalized apriori entitlement to accept intelligible assertions needs supplementation. An analogous point applies to computers' computation of theorems of hard, unsurveyable proofs.

Our question is whether the extended body of warrant of a belief in the report of a difficult theorem can ever yield apriori knowledge. More precisely, I ask whether one can ever abstract from empirical elements in the extended body of warrant and retain sufficient apriori justification in that extended body to have apriori knowledge. An affirmative answer requires that the sources' knowledge, or indicated justification, be apriori. It further requires that one's entitlement to accept the report and one's supplementary justification to maintain acceptance, despite reasonable doubts based on the theorem's difficulty, also be apriori.

We have transformed our problem into a question about the nature of the supplementary justification. Clearly, much of what we know that might supplement our basic entitlement is empirical. Knowledge of the psychology, biography, statistical reliability, or physical capacities of the mathematician or computer would be empirical. If I depended on an interlocutor that depends on such information for his or her warrant, the extended body of warrant underwriting my belief would also be empirical, even if my proprietary warrant to accept the interlocutor's assertions were apriori. Most supplemental historical knowledge of

the sort needed to lend credence to the claim of a difficult theorem will inevitably be empirical. Is there a relevant type of supplemental knowledge that might rest on apriori, non-sense-perceptual justification?

One thing to notice is that we can rely on others to provide warrant for accepting the mathematician or computer's word, without depending in the extended body of warrant *indispensably* on empirical facts about psychology, biography, or physique. A simple sort of thing that would be relevant is if several sources said that there were near proofs of the theorem, and a proof would not be a major mathematical surprise. Then one's reliance on the mathematical authority of the mathematician or computer (call her or it "Q") would be distributed through several intermediaries; and the reliance would derive from mathematical considerations that backed, for our sources, belief in the theorem. These might buttress accepting Q's word.

It remains likely, however, that in the cases of hard, uncheckable proofs, the supplemental information needed to believe what Q says would have to include something about Q's authority or competence as an individual—or at least about the authority of our interlocutors, reporting on Q. Let us assume this to be true.

Now one could be impressed not so much by Q's notoriety or physical characteristics, as by the force of Q's reasoning. One (or one's interlocutors) might have checked prior mathematical successes of the individual. The mathematical *results* are by hypothesis apriori warranted. But to support believing in the authority of an individual, the results must be attributed to *the individual, in third-person form*. Does the reliance on cumulative mathematical successes in another person or computer require empirical justification?

Recall that M had two sorts of non-empirical epistemic warrant—an induction and an entitlement. I want to discuss the third-person analogs of both sorts of warrant, for the case of relying upon Q. Let us begin with the analog to M's inductive justification.

Inductive form does not by itself make the justification empirical in either the first- or third-person case. We must consider the basis steps of the induction. Any such induction about Q has a basis in attributions to an individual in time. The induction might be of the form, "Q produced proofs p, p₁, p₂; these proofs indicate that Q is a source of great rational (mathematical) power; so Q is worthy of special epistemic reliance".

One might hold that any conclusion based on beliefs about particular events in time is *ipso facto* based on experience, and hence empirical. So any justification that bolsters Q's authority as a mathematician or computer will inevitably be empirical.

This line of thought is based on a conception of empirical experience and apriori justification that differs from mine. Frege held that an apriori truth just *is* one that has to be derived and justified from fundamental general laws rather than from assertions about particulars. Leibniz, from whom Frege got his conception of apriority, maintained that judgments about one's present thoughts (*cogito* thoughts) are judgments of experience, hence not apriori. These philosophers saw

the divide between the apriori and the empirical as lying between justifications that are founded on general rational principles, and justifications that are founded on judgments about particulars, especially events or objects in time. The latter judgments counted as judgments of empirical experience, regardless of whether the “experience” is sensory.²² On this view, neither M’s induction about herself nor our induction about Q could be apriori.

By contrast, my conception of apriori warrant features independence of justificational force from sense experience or sense-perceptual belief—not independence of any intellectual “experience”. With Kant—and the empiricisms of Mill, the positivists, and Quine—, empirical experience has come to mean sense experience. Thus apriori justificational force is force that derives from non-sense-perceptual sources—sources of reason, understanding, or perhaps the form of experience. An induction based on judgments about particular intellectual events is not *ipso facto* empirical.

It is clear that in attributing intellectual successes to another individual, we must utilize perception. The question is whether such perception must contribute to the force of our justification for relying on another individual as specially authoritative. I think that there are relevant apriori warrants that derive from intellectual understanding.

A conception of understanding that is congenial with the Leibniz-Frege conception of apriority is that intellectual understanding operates only on intentional content *types*: perception is always needed for warranting the application of understanding to particulars. I believe that this is a mistaken view of intellectual understanding. I think that intellectual understanding operates originally on token events with intentional content. One understands one’s own token judgments—the content and judgmental mode of intellectual (particular) acts. One understands them as judgements with certain conceptual content, and such content is an abstract type. One also understands the *mode or judgmental force of the particular tokening* of the type (understands its judgmental or assertive force). One cannot understand one’s own thinking or the assertions of another without grasping intellectually the judgmental or assertive force of events with intentional propositional content.

In some cases, one understands what I call intellectual *applications* of the conceptual types. For example, in understanding the first-person concept in application to oneself (or, I think, another) on a particular occasion, or in understanding some other indexical or demonstrative as applied to some intellectual act, one understands a token event. Such understanding of token events always requires conceptual guidance, but it is not reducible to the understanding of abstract conceptual (or intentional-content) types. These are cases of intellectually understanding a token event; our warrant for relying upon such understanding is non-empirical.²³

To gain conceptual understanding of one’s own thought as such, one must understand token instances. Conceptual understanding of content through language develops in tandem with understanding token judgments or assertions by

others. To understand intentional propositional content types, one must normally understand intentional content tokens that are presentations-as-true, or more narrowly, token assertions or judgments. More generally, understanding content types requires understanding content tokens. I believe that such understanding is commonly intellectual. Individuals with the relevant competence are *apriori prima facie* entitled to rely upon it. This view is embedded in our fourth assumption.

We have an *apriori* warrant to rely on our putative understanding of token assertions, or other presentations-as-true. In the first-person induction, M relied on this warrant in her self-knowledge. She was able to identify and (through intellectual memory) reidentify herself via the first-person concept. I maintained that she could be *apriori prima facie* warranted in doing so. To carry out the analogous third-person induction, one must attribute successes to the relevant source, Q. Normally we identify such a source perceptually. Our warrant for such identification is empirical. But to be *apriori* warranted in the induction that we are considering, we cannot lean on this form of identification. We have available, however, an alternative: we can fix the source as the source of the putatively understood event with intentional content—the *source or author of this presentation of intentional content*. I believe that we have an *apriori* warrant to presume that an instance of intentional, propositional content has a source with propositional powers.²⁴

In any given case, on our assumption about interlocution, we have *apriori prima facie* entitlement to accept what the source presents as true. And we (or at least some among us) can supplement this entitlement by thinking through an argument that the source gives us. If what we reason through lends credence to what the individual tells us, the individual's assertion gains credibility; and our defeasible default entitlement is supplemented with a justification. Our own going through some or all of the mathematical reasoning is by hypothesis *apriori*. So on our assumptions, *in given cases* we can obtain *apriori* justification for supplementing our *apriori* entitlement to accept the intentional contents that our source presents as true.

To yield an inductive justification for taking Q to be more than ordinarily authoritative, however, it seems, at least at first blush, that we must attribute presentations to a single source not only from the third-person point of view, but *more than once*. Must such attributions be warranted empirically?

V

In discussing an *apriori* warrant for third-person attributions, we have been implicitly advancing a position on an aspect of the problem of understanding knowledge of other minds. This issue demands much fuller development. But I want to sketch more explicitly the line that has been guiding my thinking.²⁵ Clearly, perceptual belief is a necessary enabling condition for having any warrant regarding another being's intellectual accomplishments, indeed for understanding any given accomplishment. It is a standard view that it must also enter into the

justification of any knowledge of another mind. Most would hold that one must infer such knowledge from the behavior, broadly and liberally construed, of an individual. But I do not find the matter so straightforward.

I think that one can know of the existence of another mind non-inferentially, via apriori entitlement associated with understanding intelligible assertions, or other acts with intentional content. If a reliable linguistic understander finds immediately intelligible any seemingly propositional (seeming) event, that individual has an apriori prima facie entitlement to presume that it has a rational source.

In reception from a computer (lacking careful observation!) one might mistakenly but justifiably presume that one's immediate source is a rational mind. This possibility is compatible with the relevant apriori entitlement. First, the entitlement is prima facie, and thus compatible with being mistaken. But second, the entitlement is not a warrant for the presumption that one's *immediate* source is rational. It is a warrant for presuming that there is *some* rational source in the chain of sources behind the seemingly intelligible, seemingly propositional, seeming occurrence. Although this presumption too is defeasible, it would remain not only warranted but true in the case of reception from a computer. The presumption does not depend on recognizing the nature of one's immediate source, other than as a locus and causal source of events with intentional content. The immediate source need not be a mind. Thus seeming understanding provides apriori prima facie entitlement to, and potential apriori knowledge of, the conclusion that a particular (seemingly) propositional (seeming-) event has some (ultimate) rational source, in some mind.

There are two necessary conditions on an individual's having the relevant entitlement that need to be highlighted here. The first is that the individual be a reliable linguistic understander. That is, the individual must be capable of understanding, from a third-person point of view, events with propositional intentional content according to some pattern that distinguishes, reliably, what is ("literally" or objectively) expressed from the momentary intended contents or implicatures of any one speaker. The second is that the individual must, using his linguistic capacity, find an event *immediately* (that is, non-inferentially) intelligible as an instantiation of understood propositional intentional content.

Some have wondered whether the foregoing view will justify finding rational agency behind the regularities of nature. I think not. The argument is importantly different from traditional arguments from design. The warrant attaches to the immediate application of a linguistic capacity by someone with a reliable ability to understand linguistic expression of intentional content. The physical world is not a text. It is made up of objects and properties, not (except perhaps in the case of the products of minds) intentional contents—not of entities with intentionality and some claim on truth. So perceptions of regularities in the world are normally not non-inferential applications of a linguistic capacity. In special circumstances, one might be warranted in non-inferentially taking what are in fact mere random physical events as propositional expressions. But perceptual experience of most non-propositional events does not involve application of in-

tellectual capacities associated with propositional understanding. The warrant that I am discussing attaches specifically to seeming understanding by reliable understanders of seemingly propositional events. Understanding *intentional, propositional* events—whether issuing from computers or people—lies at the basis of the apriori warrant for presuming the existence of a rational source.

Of course, my view about apriori knowledge of *other* minds depends not only on my assumptions about interlocution. It also depends on the claim that one can be apriori entitled to presume that the rational source is not oneself. I cannot defend this claim here. I sketch it only to indicate a direction of reasoning. In my view, one can have a rational, non-sense-perceptual, knowledge-yielding entitlement to believe that the source is another mind, by having non-perceptual entitlements to presumptions about one's own agency and responsibility, or lack thereof, in the reception of information.

One thinks an intentional content but takes it as received from an event that one did not produce and that has the same intentional content. That *is* third-person (putative) understanding. By recognizing that the propositional act that one is considering for evaluation is not one's own, one can rationally presume the existence of agency from *another* rational source. Again, the warrant underlying this recognition is defeasible. But I hold that the ability to recognize whether or not one is the source of understood propositional acts is fundamental to being rational. The warrant for the exercise of the ability is not grounded in sense-perception. Thus I think that one can know the existence of another mind without resting the knowledge's justificatory elements on inference from observation of behavior or on other particular sense-perceptual beliefs.

There is a sense in which the possibility of knowing apriori a computer-proved theorem does not depend on whether the rational source behind the computer's work is *another* mind. The mathematician that created the program might be the recipient of the computer's offerings. The relevance of issues about other minds is not that the computer's offerings *in fact* stem from another mind (though often they do). The relevant aspect of the parallel is that the recipient of the computer's offerings must confront them *from the third-person point of view*. Even if the computer is in some sense an amplification of the mathematician-programmer's mind, and even if the recipient is the mathematician-programmer, the recipient must consider the computer's offerings as presentations to which the recipient is not necessarily committed.

So the *epistemic orientation* of the recipient toward the rational source (the programmer-cum-computer) is the same as it would be toward another mind. We are investigating the third-person orientation toward evaluating the computer's offerings and determining the scope of its power and reliability. This orientation is like a third-person viewpoint on one's own unconscious, except that one is not in a position to access the source as one's own, in the way that M accessed her unconscious as her own. Nevertheless, it seems to me that using this third-person orientation, one can know, attribute, and evaluate the contents of the computer's

offerings without resting the knowledge's justificatory elements on inference from observation of behavior or from other sense-perceptual beliefs.²⁶

VI

Even assuming that what I have said about knowledge of other minds is correct, our problem of finding a supplementary apriori ground for believing the computer's claim to the proof of a difficult theorem is not solved. For the inductive warrant that we are considering requires an induction on the *same* source. To know inductively that some source is specially authoritative in mathematical matters, one must accumulate grounds to find the same source authoritative. So there appear to be issues about reidentification.

Many philosophers influenced by Kant, notably Strawson, have emphasized the role of perception of physical objects in reidentifying other minds, and even one's own. I value this tradition and find its basic line broadly plausible. What gives me pause is reflection on the fine line between the genetic role of perception in enabling one to learn to use a rational source and the justificatory role of perception in warranting acceptance of the deliverances of the source.

Let us return to the first-person case. I held that M could do an induction from apriori warranted basis steps, over her past proofs, to ground a conclusion that she was specially trustworthy in carrying out a new difficult argument not all of whose steps could be consciously checked. She used the first-person concept to attribute the proofs to a single source. She might have used a name whose reference and intentional content is conceptually tied to some use of *I*. But the argument was fundamentally first-personal.

Now as Kant pointed out, any use of *I* could in principle be tied to apparent memories that in fact connect not to one's own past exploits, but to those of another being.²⁷ So the use of *I* in these inductions does not *guarantee* a single source for the induction. But a guarantee is too much to ask even for an apriori warrant. The lack of a guarantee does not show that one's warrant to presume a single source requires empirical grounding. What normally provides a prima facie apriori warrant for the presumption of a single source is not a direct self-reidentification. It is, as I noted earlier, an apriori connection between uses of *I*, on one hand, and certain reliable uses of preservative memory (or more broadly, intellectual uses of memory), on the other. The ability of the individual to use intellectual memory, especially to carry on arguments over time, grounds a prima facie apriori warrant for the individual's presumption that he or she has remained constant through an induction that depends on his or her identity. Note that this ability itself carries unguaranteed but non-empirical warrants for presumptions of sameness of content, and continuity among the steps of the inference. By extension, memory of past argument steps and past intellectual acts provides defeasible, but non-empirical entitlement to a presumption, in first-person form, that one is the author of those past acts.

Is the third-person case relevantly different? We can construct a justification of the form: Q proved t_1 ; Q proved t_2 ; ...Q proved t_m ; Q has carried out a large and difficult part of the proof of t_n ; so Q is an unusually reliable source whose claim to have proved t_n can be trusted. I have maintained that if Q is identified only as the source of a given intellectual act which one presumes to understand, one can be apriori warranted in knowing and attributing any given basis step in the induction. To carry out the induction, however, one seemingly must hold Q constant. One can, of course, reidentify Q empirically. But can one justifiably presume with *non-empirical* warrant that it is the same rational source through the induction?

Before beginning to answer this question, I want to point out that it does not seem necessary that one reidentify an *individual* as rational source. It would not matter to one's warrant or one's knowledge, I think, if someone had substituted for one's computer another indistinguishable one that was made to be identical in its methods, power, and output. (It would not matter, that is, if one were not committed to its being an individual, but was rather open to there being some "source" that might take a variety of forms.) It would not matter if a committee of sources masqueraded as a single individual, but cooperated in contributing different elements in a proof. As long as the source is a single *type* of power and the individual components of the type (either members of a committee or members of a series of "replacement" individuals) are non-accidentally and rationally related, the relevant warrant and knowledge will, I think, be substantially unaffected. What is important is identity in the type of source of information, not identity of a concrete individual. I shall assume that the induction depends on reidentification of a rational source in this broad sense. But this assumption will not be critical to my argument. The points I make will, I think, carry over to defeasible attributions to an individual.

What resources do we have for non-empirically warranted re-attributions to a rational source from the third-person point of view? The role of intellectual memory in the first-person case provides a model. One might be entitled to re-attributions to the rational source through continuities in understanding the source's output, where these continuities are held together through preservative memory. That is, one uses intellectual continuities as basis for an entitlement to presume that the source is the same, rather than empirical observations to justify premises of reidentification. In turning on one's computer, one is not confined to checking the look and position of the computer. One understands what remains from the previous day as providing continuity. One goes on with the manuscript, presuming defeasibly that the resource is the same. One is not confined to reading a single screen. One can check through the document and check other documents as ways of justifying the presumption that one is dealing with the same source.

To be more concrete about the case at hand, in understanding the output of a person or computer, one has various continuities in the discourse to rely upon. One can not only continue with the proof one left off with. One can check the

person or computer's abilities and memory by having it re-prove theorems whose proving gave it special authority. One can note anaphoric connections or other memory connections both within a single argument and across arguments. One can find definitions recurring that might remain constant both within and across arguments. Such checks would, of course, be subject to error. But they could be justificational independent of sense experience, on our fourth assumption about interlocution, even though the ability to gain access to those powers requires sense experience. The relevant continuities are intellectual. The *warrant* for relying on them does not seem to depend essentially on continuities in sense perception of the computer's body. The presumption of sameness of source seems to be a justifiable *prima facie* default position, as long as one's preservative memory indicates intellectual continuities from one session and proof to another. The continuities can be relevantly like those in thinking through a discourse or in reading a novel.

Although I have represented the induction as being over a series of past successful arguments, the induction need not always take this form. What matters for establishing the special mathematical authority of one's source is not that the source prove several separate theorems or make several separate calculations. What matters is that the source somehow show relevant mathematical power and reliability. This could be shown in a fragment of a *single* argument if the fragment were complex and hard enough. In fact, there may be only one argument on which to check actual computers which have been programmed to deal with a problem as hard as the Four Color Theorem. Still, by thinking through a large portion of the proof, with its various lemmas and sub-proofs, a competent mathematician will be in a position to determine on intellectual grounds that this source has considerable power and special authority.

The fact that the induction can be carried out from the case of a single proof affects our understanding of the reidentification problem. We tend—rightly, I think—to assume that we are entitled on intellectual grounds to rely on presumptions that we are following a single argument, no matter how complex. If a source is credited with special authority through its producing a single, sufficiently difficult argument-fragment, the induction need not depend on multiple attributions at all. One need identify it only as the source of the relevant argumentation. One can induce from its virtuosity with the argument fragment that it is worthy of special credence.

Of course, what is to be counted as a specious present for these purposes is subject to possible dispute. Most arguments of the requisite complexity cannot be held in mind in one session (on any reasonable conception of "session"). This is true of the Four Color Theorem. Thinking through a substantial fragment of that proof would involve going home, sleeping, returning for more. Here, some philosophers might hold that the problem of reidentifying a source for a single argument, in these cases, is not substantially different from the problem of reidentification across different arguments.

I am inclined to think that special presumptions of continuity do attach to the thinking through of arguments (or their fragments), no matter what their (finite) length. The ability to follow and carry through arguments is so fundamental to being able to understand intentional content, and to being a single subject, that I think that this ability has a special place. But waiving this point, it seems obvious that it is in principle possible for a mathematician or computer to establish special authority through an argument that can be held in mind, by sufficiently competent recipients and with sufficient background knowledge, in a single “session”. So reidentification is not essential, in every possible case, to carrying out inductions that support a source’s special authority.

Even if one lays aside appeals to single sessions, the continuities in a single argument—no matter how long—provide a defeasible intellectual basis for presuming a single source. As I have indicated, there are further intellectual continuities that hold regardless of whether one is dealing with a single argument. These continuities seem to me sufficient to solve, on intellectual grounds, the problem of reidentification for purposes of carrying out a successful induction—regardless of whether the induction takes as its basis a single intellectual achievement or a series of them.

So far I have explored how giving the source special credence in its claims to have proved a difficult theorem rests on an apriori-based induction on the source’s virtuosity in solving mathematically difficult problems. But there is a dimension of difficulty different from mathematical difficulty that I have not discussed—namely, sheer length. One might think that a substantial element in possible doubt as to whether the machine has proved the Four Color Theorem is the length of the proof. One seems to need a warrant to believe that the machine is “big enough” to have carried out the proof. And this warrant might seem to be obtainable only from empirical information about the storage capacity of the machine. This (one might think) could not be gained from induction on past proofs that had been checked. For by hypothesis, this proof is too long to check. The whole issue might seem to be whether the machine has the storage capacity to deal with a proof that requires x bits of storage capacity, where x is much larger than the length of any checked proof. How can one be *non-empirically* warranted in presuming that the machine has carried out a proof that is longer than any proof that one could think through?

I believe that this problem is less difficult than it may first appear. How do we know how long the proof in fact is? The machine is (or might be) equipped with a counter that counts steps in the proof. By understanding the content of the output of the machine, one could know inductively that the counter is reliable in counting steps. So an induction that projects beyond the actual lengths one has checked could be warranted. Such an induction would proceed from mathematical, apriori basis steps. Similarly, one could have recognized that the machine announces a proof (stops and produces the turnstile) only when it has one—to an extremely high degree of reliability. Its announcing a proof in the

case of a long proof could be deemed credible on inductive grounds. Again, the basis steps for the induction seem to be warranted non-empirically. We can recognize assertion of theorems through understanding. This is a corollary of our fourth assumption. And the recognition of the mathematical probity of the individual proof-, or sub-proof-, announcements is mathematical, hence apriori by our first assumption. The reidentification problem is substantially the same in cases of determining mathematical endurance as it is in the cases of determining mathematical virtuosity.

In summary, one can be apriori warranted in a third-person inductive argument of the form: Q (identified as the rational source of the proofs) has carried out the proofs $p_1, p_2 \dots p_m$; these provings indicate that Q is a source of great mathematical power; Q claims a proof p_n and has produced a large relevant fragment of such a proof; so Q can be given special credence in Q's claim to have proved p_n . This is how the apriori warrant goes: One is apriori entitled to one's understanding of a sequence of contents presented as a proof. One can presume apriori that that sequence has a rational source. Call the source "Q". One can be apriori warranted in believing that the sequence in fact constitutes a proof by thinking through the sequence and realizing that it is a proof. Call the sequence "p1". Using intellectual memory to hold the initial understanding together with the mathematical replication, one is apriori warranted in believing that Q has carried out the proof p_1 . Similarly, for $p_2 \dots p_m$, and the fragment of p_n . One can hold this sequence of proofs together by intellectual memory. And one can be warranted in thinking that the proofs have a single source, Q, by noting intellectual continuities among the proofs, and by rechecking Q's ability to carry out proofs again. One is in a position to appreciate, on mathematical grounds, the special mathematical ability of Q as indicated in these proofs by appreciating the mathematical difficulty of finding and carrying out the proofs. One's use of the inductive form of argument, as well as one's beliefs in the inductive base (just sketched), are apriori warranted. So the whole induction is apriori warranted.

I conclude that, given our assumptions, one can be in a position, from the third-person point of view, to be apriori warranted in believing, in fact knowing, on defeasible, inductive grounds, that the theorem has been proved. One can know this even if one cannot fully replicate the proof. The problem of reidentification in relevant inductions can in principle be solved through apriori warrants. I do not presume to have discussed the reidentification problem in appropriate depth. But I think that I have sketched a view that is plausible and deserves development.

VII

We have been discussing third-person, apriori, inductive warrants for granting special authority to an interlocutor who asserts a difficult theorem. Recall that M, our first-person reasoner, had two forms of non-empirical warrant—an induc-

tive form and an entitlement that does not rest for its warrant on an induction. I want to turn to the third-person analog of the non-inductive entitlement.

I took M to be *non-inductively and non-empirically entitled* to rely on her mathematical powers because they are in fact rational and she knows how to use them with sufficient understanding. Is there a third-person, non-inductive, apriori entitlement to accept the offerings of some other person or computer, Q, as specially authoritative?

Through use of her unconscious mathematical powers and through understanding their methods and products, M acquired an entitlement to rely on those powers as more than normally powerful. In being an entitlement, the warrant does not involve a meta-induction on past successes. The warrant is the counterpart, for a rational faculty, of an individual's entitlement to rely on the sharpness of a very sharp perceptual faculty, without doing a sophisticated induction on its comparative merit. I want to consider whether there is a third-person analog of M's reliance on her own unconscious mathematical powers. Such an analog would be an entitlement to rely on Q as a powerful amplification of one's own rational faculties. The delicate matter is to articulate an apriori entitlement that meets two conditions. It must reside in knowing how to use a cognitive faculty or source, not in an induction on products of the source. But it must be in third-person form—an entitlement to rely on the products of a rational resource that is not taken to be one's own.

For an individual that has the appropriate relation to Q, the entitlement would take the form: *the powers one relies upon (Q's) are rational and are more than normally powerful, and one knows how to use those powers and understands their methods and products to a degree sufficient to appreciate their power.* I think that there are two primary aspects of *knowing how* to use Q's powers in a way that entitles an individual to rely on those powers more than one might rely on those of any arbitrary rational being.

The first aspect is a kind of "know how" that is not to be distinguished from tracking some of Q's intellection—thinking through reasoning that one correctly (and with warrant) attributes to Q. For example, one might follow some of Q's proofs or sub-proofs. One need not have any idea how Q targets potential theorems, or how Q discovers the proofs. It is enough that one think through Q's justifications for theorems, understanding them well enough to recognize their soundness and appreciate the power and difficulty involved in coming up with justifications for the relevant results. This appreciation need not be meta. One need not think *about* a series of past proofs. Nor need it involve comparisons between Q's intellection and that of other mathematicians whose accomplishments one thinks about from the third-person standpoint. It is enough that one understand the justifications sufficiently to realize *through understanding them* that they involve special powers (perhaps more power than one has), and sufficiently to understand their ability to solve problems whose difficulty one appreciates. One gains an appreciation of Q's power by being expanded by it, or at least by internalizing fragments of it.

To this degree, one's appreciation of Q's power seems similar to a great mathematician's appreciation of his or her own power. Such a mathematician *might* do an induction on his or her past successes or a comparison with the abilities of others. But to be entitled to rely on the power, the mathematician need not take such a meta point of view. The mathematician's basic entitlement derives from using the power in a way that enables appreciation "from the inside" of the type and difficulty of the problems that he or she can solve.

For contrast, consider a case in which one has no insight into the reasoning of one's source. One simply finds that he, she, or it gets answers to difficult problems that when checked are found to be sound. Assuming that one has some ground to recognize the other's methods as rational, one can be warranted in trusting the other's assertions to a greater than normal degree. A historical approximation to this sort of case is Hardy's increasingly warranted confidence in Ramanujan's mathematical results, despite at best sketchy justifications of them by Ramanujan.

I believe that one can have a non-empirical *inductive* warrant, for relying on such a source to a greater than normal degree. But it seems more questionable that in such cases, where one has no insight into how one's source solves its problems, one can have a *non-inductive* apriori entitlement to rely on the exceptional rational powers of one's source. If one cannot incorporate any part of the source's ratiocination, there is no clear sense in which one is using the other's powers as an amplification of one's own. One is not in a position to appreciate one's source's powers of inferential justification "from the inside", by reproducing them in one's own cognitive processes. If one can replicate none of the source's inferences, one does not know how to use the power—except in the weak sense that one can recognize, check, and come to rely on its *results*. I doubt that this weak type of know-how suffices to provide a non-inductive entitlement in this case. It seems to me that in such a case one is forced to treat the source as an object whose special rational reliability can be established only through induction on its past successes. The non-inductive entitlement depends, I think, on an ability to "incorporate" the ability (at least partly) into one's own point of view. So I will assume that the first aspect of the relevant sort of know-how with *mathematical* powers involves being able to follow some of the justifying reasoning, and to appreciate, from the reasoning itself, its degree of power and insight, as well as its reliability.²⁸

Of course, insofar as this first aspect of knowing how to use the source's rational powers involves simply mathematical reasoning, it contributes, given our first assumption, no empirical element to the entitlement to rely on the source as specially authoritative. The question of whether the putative non-inductive warrant can be apriori hinges on a *second aspect of knowing how to use another's* rational power in some domain. This second aspect concerns knowing how to *access* the power, indeed how to access the same power (or relevantly type-identical power).

With first-person appreciation of a mathematical power as special, one's access to the power's products can be through self-knowledge. I maintained that

this access is apriori warranted. But in accessing the power of another source, one does not have first-person access. According to our fourth assumption, one can have apriori warrant for one's understanding of what another being says. But can one have third-person apriori warrant for a reliance that depends on *re-using* the same power over time in a way that does not involve meta-attributions of the results to the power, yet yields a non-inductive, third-person apriori entitlement to the use of that power?

Let us consider another parable. Suppose that to access mathematical reasoning in oneself, one had to see a red light. Suppose that to access higher mathematical reasoning, say analysis, one had to see a red light of a certain type juxtaposed with a green one of another type. Suppose that one is very good at analysis. As one solves more difficult problems, one develops, through the practice of one's reasoning, a warranted confidence in one's ability to solve difficult problems. Part of one's competence lies in an ability to determine whether a given problem is within the range of one's powers. To access and develop one's competence in analysis, one would have to use one's senses to reidentify the types and colors of light. But such use would clearly not, I think, be a part of a basic warrant for relying on one's reasoning or one's special abilities in analysis. It would merely be part of knowing how to start the reasoning on any given occasion. Empirical cues that go into accessing one's rational powers are not part of one's entitlement for relying on those powers to whatever degree one is entitled to rely on them. They are part of knowing how to use them.

One could, of course, construct an empirical induction: because one had gotten reliable answers to difficult problems after seeing the relevant lights, one could rely on one's answers to difficult problems in analysis after seeing the lights. But such an induction would, I think, be an unnecessary superimposition on a non-inductive entitlement that would reside in knowing how to access and use one's mathematical powers.

In such a case, one would be entitled to more-than-ordinary reliance on those powers because they are rational and more than ordinarily powerful, and because one knows how to use them, in the sense that one understands their methods and products to a degree sufficient to appreciate their power. The fact that one has to see certain types of lights to use one's powers bears only on one's knowing how to access the powers, from the standpoint of this entitlement. The perceptual experience, though a necessary condition for using one's mathematical powers, does not contribute to the force of the warrant, however psychologically essential it is to the exercise of one's powers. The warrant is a non-inductive apriori entitlement.

Learning a language and learning certain types of mathematical reasoning both require use of one's senses. One must hear the words, see the symbols, reflect on the diagrams, to acquire the abilities to reason mathematically. But such reliance on perceptual belief is normally not part of one's warrant for accepting the propositions that one comes to understand in this way. Similarly, in accessing and coming to exercise one's mathematical powers on any given oc-

casation, perceptual belief could have been necessary. But even if it were, it need not function justificationaly in one's non-inductive entitlement to rely on the rational capacity to a degree appropriate to its rational power.

In the third-person case, it is not obvious to me why perception of a computer, like perception of its symbols, cannot be a cue for drawing one's attention to a resource for one's rationality, rather as seeing red and green lights might be a pre-requisite for using one's higher mathematical powers, but no part of warranting belief in their offerings. The perceptual experiences needed to tap into the resource would be part of *knowing how* to gain entry to the resource, but no indispensable part of a warrant for relying on it to whatever degree one's actual mathematical use and understanding of it allows. Using Q as a rational resource would require empirical recognition. But the recognition would only be a matter of knowing how to access the source, from the standpoint of a non-inductive entitlement to rely on its mathematical power. It would no more contribute to the force of the relevant non-inductive warrant for relying on the source than seeing the lights would contribute to the force of a warrant for relying on one's own rational powers. The entitlement would derive from one's employment and understanding of products and procedures of the mathematical resource, given that the resource was accessed.

Once one has accessed the source, an appreciation of its exceptional mathematical powers can be gained, as we have seen, by reasoning through its solutions to difficult problems (even perhaps *one* difficult problem), and understanding "from the inside" the level of problem that the power can deal with. The warrant for resting more-than-ordinary confidence in the source need not be inductively based. One's warrant resides in one's understanding the source's solutions to difficult problems. The warrant resides in one's ability to reason through, and appreciate the reliability of, the source's reasoning with respect to a range of difficult problems. By incorporating Q's mathematical work, and understanding it, one's warrant to rely on the source expands to fit the level of power that the source has exhibited.

The third-person, non-inductive apriori entitlement is analogous to the non-inductive entitlements we have discussed before. *One is entitled to more than ordinary reliance on Q's powers because they are rational and are more than normally powerful, and because one knows how to use those powers and understands their methods and products to a degree sufficient to appreciate their power.*

What happens to the problem of reidentification, with respect to the non-inductive entitlement to rest more-than-ordinary reliance in this source? Part of what it is to be entitled to rely on Q is not to need to justify the reliance in terms of attributions to Q, of the sort that formed the induction base of the inductive justification for relying on Q. So to be apriori entitled to such reliance, the recipient need not, in beliefs about Q and Q's offerings, *refer to* Q and Q's offerings as objects. The recipient need not refer to Q even in the form "the source of these [intellectually understood] presentations". The recipient must access Q's inten-

tional contents through perceiving Q's symbols, and perhaps Q as well, to trigger intellectual understanding of the products of a single, coherent resource for rationality. But the "God's eye" articulation of the recipient's entitlement to rely repeatedly on Q can specify Q only as the source of the offerings that Q understands. And among these offerings, there is sufficient continuity, of the sort discussed in connection with the inductive justification, to entitle the recipient to presume on the sameness of the source.

No analog of the reidentification problem arises, of course, if one's appreciation of exceptional intellectual power derives from one session with Q (thinking through a large sub-part of a single proof, or, given enough background knowledge, appreciating the power and difficulty of a single, short, creative justification). But the single-session case is not essential to the point. One's warrant can be non-inductive in the normal multi-session case. The "God's eye" articulation of the entitlement need not refer to the empirical characteristics of the source. It can specify Q as the source of a series of events with intentional content that the recipient understands. Where there are appropriate continuities among the intentional contents and proofs (continuities of anaphora, definition, level of ability, and so on), the recipient can be entitled to rely on the sameness of the source even if the recipient cannot conceptualize the continuities as such.

M needed to *know how* to access and use her powers. But to have warranted confidence in them, she did not need to reason about them, beyond understanding and checking their products. The account of the warrant to rely on one's own reasoning clearly does not require that one reidentify one's reasoning capacity in that domain whenever one accesses it. Even if one must see lights in order to access one's capacity, those perceptual pre-conditions need not be specified in an induction that depends on one's reidentifying one's reasoning capacity each time one accesses it. The capacity need not figure as an object in one's reasoning. Nor need the recipient's dependence on perception in accessing the source be specified in the non-inductive entitlement which explains the recipient's rational right to rely repeatedly on the source's reasoning power. It is enough that there be actual identity of the power and perhaps some simple continuities among its products. One's ability to regularly reaccess the capacity suffices to enable one to be entitled to rely upon it.

I am hypothesizing that certain third-person cases can be similar in these respects. In reaccessing the source, one perceives its symbols and its empirical characteristics; but they play the role of triggers for access, for the immediate application of intellectual understanding, and for reliance on the source's intellectual continuities. But one need not *reidentify* it as part of an induction. One need not refer to it or its products as *objects* in one's reasoning at all, anymore than one need refer to oneself or one's own thoughts in being entitled to carry out one's own reasoning. One learns how to use the source's rational power and to rely upon it through understanding its reasoning. The entitlement is based on the

rationality of the resource and on one's understanding its activity and knowing how to use it as an extension of one's point of view.

Despite the fact that, as recipient, one need not reidentify or refer to the source (or its intentional content) as an object in justificatory reasoning, one does bear an objective, third-person relation to the source. What does this third-person relation consist in? I think that there are three relevant features. First, there is a rational gap between its asserting something, or presenting it as true, and one's own acceptance of the same thing. Its presentations are not part of one's point of view unless one incorporates them. Second, there can be failure of access and failure of (non-explicational) understanding—failure of comprehension—without there being any rational failure or malfunction in one's own cognitive system. For example, one can mistakenly presume, with warrant, that a putative source is offering propositional contents, when there is really only random noise. Or one can access a source that appears to be the same as the source last accessed, but it is not. These failures need not signal any malfunction or rational failure in one's own point of view. One is subject to *brute error* with respect to one's presumptions. Third, any rational being with appropriate abilities of understanding could have been equally well placed to understand and make use of the relevant source. And such understanding could have had the same epistemic status or warrant with respect to that source.²⁹

But the *objective*, third-person relation to the source does not entail that the individual recipient need make reference to the source as *object* in some justification or warrant. The recipient must perceive the source and expressions of its intentional content. But this perceptual identification need not be specified in the account of the individual's entitlement. The entitlement for relying on the source does specify the source. But it need not specify the object's empirically discernible characteristics, or the empirical background conditions that enable the recipient to access and rely on the source. It can specify the source in the non-empirical way that the apriori inductive justification does. With respect to this non-inductive entitlement, perceptual access is as much a matter of know-how and as little a matter of object reidentification by the recipient as it is in the first-person red-light green-light cases.

It is a delicate and unresolved matter (in my mind) how to distinguish the cases in which *warrant* for continuing reliance on a source Q requires an empirical induction, or even an empirical entitlement, from the cases in which empirical recognition can be submerged into knowing how to access a rational resource. The latter cases seem to me to require that the perceivable properties of a computer or person that one uses as rational resource be relatively simple. I think that they must be incorporated into a nearly automatic routine. It is important that the recipient need not engage in context-dependent empirical (or non-empirical) tracking exercises, or complex theorizing, to reidentify the resource (type or token) through its possibly changing physical characteristics. We rely on empirical constancy for access, without having to refer to it in our warrant.³⁰

If the world were such that we did have to engage in complex theorizing to access and reidentify a rational source, then perhaps we could not avoid relying on empirical reasoning to justify a presumption of repeated reliance on the authority of the same person or computer. Access-know-how would not simply be the occasion for triggering intellectual understanding. The application of the understanding would require an empirical justification.³¹

As things stand, the perceptual properties of a source seem capable of functioning as triggering codes for accessing rational processes, rather as familiar words are ways of accessing meaning. Perceptual experience of the words or of the body of the source need play no role in justifying one's understanding of, or intellectual uses of, the content of the words or the presentations of the source. Perception is submerged into know-how: the conditions that make possible the intellectual understanding and continued use of a resource for rationality. Perception is only the mode of access, an enabling condition which makes no contribution to the epistemic force of the warrant.

The distinction here between background enabling conditions and justificational force is fundamental to understanding the contributions of reason and sense-perceptual experience in providing warrant for our beliefs. What is difficult is separating the role of perceptual reidentification in justifying an inference from its role in enabling one to know how to use a resource for rationality on which one is entitled to rely. Although I know of no recipe for applying it, the distinction is fundamental to developing a reasonable and nuanced epistemology. Such an epistemology must recognize the ways in which reason and understanding underwrite many of our cognitive practices, even though being able to reason and understand depend—both in general and often in particular cases—on an enabling background of sense-experience. I hope that the foregoing discussion will contribute to a deeper understanding of the distinction.

I conjecture that if perception is used only to access the mathematical offerings from a single source which are then checked for validity or plausibility and found to be more-than-ordinarily reliable and powerful, then the perceptual features of the object need not play a role in one's inductive justification or non-inductive entitlement for relying on it. Insofar as one's reliance on the computer as a powerful mathematical source is the product simply of an association of perceptions of a given look, name, or model tag with a series of powerful pieces of understood mathematical reasoning, I conjecture that perception is merely a non-justificatory pre-requisite to special reliance on an authoritative source for intellection. Perception is merely a condition that enables one to make use of a resource for reason and understanding. I believe that one can, in many cases, treat recognitional aspects of another person or computer as keys to tapping a resource for rationality, in the way that one might use empirical triggers for accessing or starting-up one's own rational faculties in some domain. I am less sure of my grip on the apriori non-inductive entitlement that I have discussed than on my grip on the apriori inductive justification. But I think that both are types of apriori warrant for accepting difficult mathematical results from another being.

VIII

As one comes to learn from the computer, to understand and rely on its arguments and results, it becomes analogous to one of one's own rational faculties. One's empirical activity comes to be submerged into one's knowing how to use the computer as an extension of one's own rationality. At least for those who understand the mathematics and rely on other mathematically adept persons or on computers, the knowledge that they obtain from understanding, checking, and appreciating the relevant mathematical work seems sometimes to rest on entitlements or justifications underwriting use of resources for reason. For those people, perception appears to be playing only an enabling, not a justificatory role with respect to some of their warrants. They will normally also have empirical justifications for their reliance on the computer or person. But their knowledge that the theorem has been proved can be underwritten by an extended body of warrant that is apriori. In these cases, other mathematicians, or computers, can be sources for apriori justification of one's knowledge of the source's result. Apriori knowledge can be transmitted through communication—even when the recipient cannot alone justify his knowledge, and even when the source must be accorded special authority if reliance on it is to be warranted.

I offer these reflections on the line between apriori justification or entitlement and non-justificatory empirical background conditions in a conjectural spirit. There are numerous difficult issues that I have not gone into. I commend the subject to reflection.

Notes

*I want to thank Joseph Almog, Tony Anderson, John Bartholdi, Jon Barwise, John Carriero, David Kaplan, Bill Hart, Paul Hovda, Tony Martin, Chris Peacocke, Tony Roy, Jim Tomberlin, the UCLA Philosophy of Mathematics workshop—which scrutinized a presentation of the paper over six all-afternoon sessions—, and audiences which heard versions of the paper at various times from 1991 to 1997 at CUNY, MIT, Oxford University, Santiago de Compostela, University of Southern California, University of Barcelona, University of California at San Diego, University of California at Northridge, and University of Illinois at Chicago.

1. There are also cases in which computers provide overwhelming inductive evidence for a conclusion. I think that much of what I say will carry over to these cases, but I will not assume they are similar in every way. Cf. M.O. Rabin, "Probabilistic Algorithms" in *Algorithms and Complexity: New Directions and Recent Results*, J.F. Traub, ed., (Academic Press, New York, 1976).
2. Saul Kripke, *Naming and Necessity* (Cambridge, Mass., Harvard University Press, 1980), p. 35.
3. Thomas Tymoczko, "The Four-Color Problem and Its Philosophical Significance" *The Journal of Philosophy* LXXVI (1979), p. 63.
4. I mention three cases. Leibniz emphasizes the point repeatedly. Cf. e.g. at the very beginning of *New Essays on Human Understanding*, Preface, pp. 48-9. Kant makes

the point at the beginning of both editions of *Critique of Pure Reason*, A1, B1. Frege writes, “If we call a proposition empirical on the ground that we must have made observations in order to have become conscious of its content, then we are not using the word “empirical” in the sense in which it is opposed to “a priori”. We are making a psychological statement, which concerns solely the content of the proposition; the question of its truth is not touched.” *Foundations of Arithmetic*, J.L. Austin trans. (Northwestern University Press, Evanston, Illinois, 1968), section 8. For me, it is not so much truth as warrant that is at issue. I will speak of the *warrants for believing* propositions (intentional propositional contents), rather than the propositions themselves, as apriori or empirical. But the fundamental point is the same.

5. Although justifications or entitlements are what are primarily apriori, I count a justificational method or resource apriori if someone who used that method or had that resource could have apriori justification or entitlement through it, or if our entitlement to rely upon it does not rely for its force essentially on sense experience.

A different application of “apriori” is to truths. I am sympathetic with this usage. I leave as an open problem the relation between my application (to warrants, methods, capacities, uses of concepts) and this application (to truths). I am inclined to think that the apriority of truths is to be explained in terms of the apriority of possible, canonical warrants for believing those truths. But the issues are complex.

6. Cf. my “Content Preservation” *Philosophical Review* 102 (1993), pp. 457-488, for more detailed exposition of this conception, as well as for some defense of the fourth assumption set out below. I shall say more to support this assumption in the course of this paper.
7. As I will indicate in discussing my third assumption, ordinary mathematical reasoning by working mathematicians often involves taking for granted theorems passed on by other mathematicians, even if the recipient does not know the proof. Since I think that it is plausible that the mathematical enterprise as a whole, in its communal form, is apriori warranted—grounded in intellection—I believe that there is some force to arguing from the dependence of mathematicians on one another for their knowledge, to the possibility of preserving apriori knowledge through communication—that is, to the truth of my fourth assumption below. In other words, since mathematical practice is apriori warranted, and since mathematical practice depends on knowledge being preserved through communication, individuals can and do obtain apriori knowledge through communication (even without going through all the reasoning that underwrites the knowledge). I will not argue in this way here. I am fully aware that philosophers with a Cartesian bent will not find this route to my fourth assumption persuasive. Moreover, I think that the fourth assumption can be argued for independently of assumptions about mathematical knowledge. Since the present discussion will use these assumptions as starting points, however, I will not here argue for them directly at all. To divide the labor among my assumptions as cleanly as possible, I will apply this first assumption, in the course of this discussion, only to mathematical knowledge for which the *individual’s* own thinking and understanding provides sufficient warrant for his or her knowledge. So for the sake of my argument, the assumption is to be taken to mean that when an individual obtains knowledge of propositions of pure mathematics through his or her own mathematical understanding or reasoning, the knowledge is ordinarily apriori. Thus insofar as the warrant for the reasoning and understanding is the individual’s own, and is specifically mathematical, the warrant is apriori.

8. "Content Preservation", *op. cit.*; and "Interlocution, Perception, and Memory" *Philosophical Studies* 86 (1997), pp. 21-47.
9. I am oversimplifying. The conditions of dependence on previous members in the chain are complex. One's immediate interlocutor might perhaps not be knowledgeable. But the interlocutor must either pass on knowledge that reposes somewhere in the chain or, as a computer might, otherwise indicate the existence of grounds for knowledge.

There are other sources of Gettier cases in which the recipients's true, warranted belief fails to constitute knowledge—beyond those that depend on lack of knowledgeability in the chain of sources. For example, a recipient may be told something true by a knowledgeable source, may believe it, and may be warranted in believing it. But conditions in the context may be such that if the recipient knew them, he would lose his warrant to believe the interlocutor. To give a slightly fanciful but provocative example, suppose that Fermat had proved (hence known) his last theorem along the lines that it was actually proved; and suppose that a contemporaneous, fellow mathematician, knowing Fermat's general reliability as a mathematician, and thinking reasonably that the proof was relatively easy, had believed Fermat. Then the recipient would have been warranted, let us suppose, in believing Fermat. If the theorem had been no harder than some ordinary equation in three variables, perhaps the recipient could have obtained knowledge of the theorem through the interlocution. But in view of the fact that the proof is vastly harder than anyone had realized (other than Fermat, on our supposition), it seems plausible that the mathematician's warranted true belief would not have constituted knowledge—even if Fermat had produced the proof in the privacy of his study. If the mathematician had known how hard the proof was, he would have needed supplementary reasons to be warranted in believing Fermat's assertion—reasons beyond knowing that Fermat was a good enough mathematician to have solved the ordinarily difficult problems of his day. For Fermat's proof of Fermat's last theorem would have required substantially more mathematics than any of the other problems that Fermat had solved.

The problem of Fermat's last theorem is notorious for having yielded false reports of proofs by excellent mathematicians. In fact, of course, the eventual prover, Michael Wiles, first reported erroneously that *he* had proved it. Within a couple of months an error was found. It was fifteen months from the time Wiles made his highly publicized announcement to the time when he closed the gap in his erroneous proof and produced a correct one. It is widely assumed that the proof is much too hard for Fermat to have discovered.

10. The recipient might hear a statement from a source who is not knowledgeable, but find the statement self-evident, or otherwise mathematically convincing, or might immediately see a proof, or might find justification in the fact that the statement coheres with other things he knows. But in such cases the recipient is relying not on interlocution but on other resources in his or her proprietary justification for knowledge. The individual's warrant for the knowledge does not reside in the interlocution.
11. Tymoczko, *op. cit.*, p. 72.
12. Tymoczko holds that neither the argument run through by the computer, nor the argument given by the programmers, is a proof in the traditional sense, because both are unsurveyable. He thinks that counting either argument a proof would be to change the traditional meaning of "proof". I agree with Paul Teller in "Computer Proof" *The*

Journal of Philosophy LXXVII (1980), pp. 797-803, that the computer has completed a proof in the traditional sense, and that the requirement that a proof be surveyable by humans is not part of the sense of the word “proof”. My argument does not, however, hinge on the outcome of this dispute. Teller sees Tymoczko as arguing directly from the unsurveyability of the computer’s work to the empirical character of the human argument. There is some suggestion of this argument—for example, Tymoczko, *op. cit.*, p. 74. But I am not sure that this is an accurate interpretation.

13. It is a delicate question whether to count the computer programmers’ argument a proof—an elliptical proof that involves appeal to the computer runs—or as merely a non-demonstrative argument that the computer has completed a proof. I think, however, that the answer to this question does not matter for our purposes.
14. Lakatos has highlighted the variety of types of non-demonstrative reasoning in mathematics. Cf. Lakatos, “What Does a Mathematical Proof Prove?”, and other papers in *Mathematics, Science and Epistemology*, Worrall and Currie eds. (Cambridge University Press, Cambridge, 1978). Cf. also George Polya, *Induction and Analogy in Mathematics* (Princeton University Press, Princeton, 1954). As I have been emphasizing, the reliance on the calculus before it received its foundational explications in the 19th and 20th centuries seems to have required non-demonstrative reasoning. But I think that competent mathematicians had apriori knowledge of the calculus. The basic principles were not (and were not seen as) axioms, or as theorems derived demonstratively from axioms. It seems to me very implausible to think that the role of the calculus in physical theory was the only sufficient ground for accepting it as a part of mathematics, or as true, before the calculus received its modern formulations. Mathematical fruitfulness seems to have been the most significant ground. My suggestion is that the ability to see fairly deeply into an unsurveyable proof might give one mathematical but non-demonstrative grounds for believing its conclusion, especially given that one has reason to think that the computer has carried through the full proof. The epistemic status of this reason still needs discussion. But the fact that the relevant proof is unsurveyable, and the fact that the warrant that a mathematician has is non-demonstrative, clearly do not *themselves* force the mathematician’s warrant to be empirical.
15. Michael Detlefsen and Mark Luker, “The Four-Color Theorem and Mathematical Proof” *The Journal of Philosophy* LXXVII (1980), pp. 803-820, think that the fact that ordinary mathematical proofs involve calculation that may fail to correctly realize a mathematically sound algorithm by itself shows that proofs are empirically justified. This view seems to me quite mistaken. I think that it involves one or more of the following conflations: empirical reasoning with fallible reasoning; aids that are relevant to, or necessary to, the psychology of proof with aspects constitutive of the justification associated with giving or accepting the proof; meta-knowledge about the circumstances in which proofs are given with the reasons actually given in a proof; procedures for checking against possible doubts with reasoning procedures that may justifiably be relied upon in the absence of reasonable doubt. I shall not pursue these issues here. But I think that this sort of empiricism about mathematical epistemology shows the effects of decades of philosophers’ not taking rationalism seriously enough to understand it.
16. The notion of autonomous rationality has been a matter of dispute. I take it seriously. Of course nature “fashioned” us. But nature does not have a mind.

17. I am relying here on the fact that the programmers of the computer that proved the Four-Color Theorem had considerable insight into the specifics of the proof strategy that the computer carried out, even though they could not go through all the details of the proof. The proof was what is known as a “computer assisted” proof. There are cases in which one can know that a proof has been carried out but have virtually no insight into how. There has recently grown up an area of mathematics which studies the properties of what are known as “Zero-Knowledge” proofs. Cf. Goldwasser, Micali, and Rackoff, “The Knowledge Complexity of Interactive proof Systems” in *SIAM Journal of Computing* 18 (1989), pp.186-208; Goldreich, Micali, and Wigderson, “Proofs that Yield Nothing But Their Validity or All Languages in NP Have Zero-Knowledge Proof Systems” *Journal of the Association for Computing Machinery* 38 (1991), pp. 691-729; Blum, Santis, Micali, and Persiano, “Non-Interactive Zero Knowledge” (1990) reprint from MIT Laboratory for Computer Science. I leave open wherein the epistemology of zero-knowledge proofs is different.
18. I am assuming that these mathematicians can know the proof strategies of their unconscious calculations because they can bring to consciousness fragments of those proofs as they are being carried through, just as many human fast calculators have some insight into the methods of their mostly unconscious calculations. So the hypothesized mathematicians can know that the conscious checks are checks of (parts of) the same proofs that they had carried out unconsciously. We could imagine cases in which the mathematicians lacked this insight into their unconscious methods. Then their conscious checks would be checks only of the results of the unconscious calculations. These checks would still lend credence to the unconscious calculations, even if it were not known what sort of unconscious proof had been carried out, or even whether the result had been established through a proof.
19. This is of course a complex and disputed issue. It is the subject of Kant’s third paralogism, and Kant seems to take a position opposed to mine. I do not hope to do justice to this matter here, but I shall discuss the matter somewhat further below—cf. note 27.
20. See “Content Preservation”, *op. cit.*; “Our Entitlement to Self-Knowledge”, *Proceedings of the Aristotelian Society* 1996, pp. 1-26; and “Interlocution, Perception, and Memory” *op. cit.*, for discussions of the apriority of self-knowledge and preservative memory.
21. “Content Preservation”, *op. cit.*; “Interlocution, Perception, and Memory” *op. cit.* Tymoczko *op.cit.* and Detlefsen and Luker *op. cit.* place heavy emphasis on the view that knowing that the computer has given certain results is empirical, as does Michael Resnick, “Computation and Mathematical Empiricism” *Philosophical Topics* (XVII) (1989), pp. 129-144. Much of Resnick’s paper is devoted to defending the modest (and I think correct) conclusion that sound argumentation for mathematical conclusions *can* involve empirical premises (e.g. premises about the empirically known reliability of a computer or human). See especially *ibid*, pp. 133-137. But Resnick infers from his arguments that these empirical methods are for us “largely ineliminable”—presumably ineliminable from the warrant for the mathematical results (p. 141). I think that he has not established this stronger conclusion or even seriously argued for it.
22. Gottlob Frege, *Foundations of Arithmetic*, section 3; G.W. Leibniz, *New Essays*, Book IV, Chapter IX.
23. For more on applications and intellectual understanding of them, see my “Belief *De Re*” *The Journal of Philosophy* 74 (1977), pp. 338-363; and “Interlocution, Perception, Memory”, *op. cit.*, e.g. notes 3 and 12.

24. The powers can be either autonomous or derivative for these purposes.
25. I want to emphasize immediately that what I am about to say is not meant as an answer to scepticism about other minds, although it may play a role in such an answer. I intend to describe one of our actual warrants for our knowledge of other minds. But I do not discuss what the legitimate sceptical questions are and whether the approach I am taking avoids begging such questions. These are further tasks.
26. Given the conceptual tie between intentional events and there being an ultimate rational source of those events, and given an apriori reliable capacity for determining whether the events are one's own acts or another's, one can be apriori defeasibly warranted in presuming that the ultimate source of the relevant intentional acts is another mind. My point is just that in order to know the computer's offerings, one need not *know* whether the ultimate source is *another* mind. One could even be warranted in believing that it was, and be mistaken. The induction that we are discussing—the one needed to lend credence to the computer's special competence in presentation of a difficult theorem—can be carried out with a third-person attribution of relevant intellectual accomplishments, without taking a position on whether the source has autonomous or derivative rationality, or whether the ultimate rational (autonomous) source is oneself or another. Cf. note 24.
27. Immanuel Kant, *Critique of Pure Reason* A361-5. The notion of quasi-memory was later introduced by Sydney Shoemaker, "Persons and Their Pasts" *American Philosophical Quarterly* 7 (1970), pp. 269-85, reprinted in Shoemaker, *Identity, Cause, and Mind* (Cambridge, Cambridge University Press, 1984).
28. The issue about whether one could have a non-inductive non-empirical warrant for relying on another being as a rational source of mathematical information in the absence of insight into the source's justifications bears comparison to the issue over whether one could be entitled non-inductively to rely on a capacity in oneself to have correct answers to difficult mathematical problems simply pop into one's mind, without the slightest understanding of why they are true, or of how to use these results in further reasoning. One could learn, by checking through normal inferential means, that the answers one gave were very reliable. One could support one's acceptance of future answers through induction on past successes. But I doubt that one could be non-inductively entitled to rely on this manifestly inferential, non-perceptual capacity without justification simply because it *was* reliable. I am supposing that one lacks the slightest insight into why the answers were true, how one got them, or how to use them further. I believe that an individual's free-standing (non-interlocutionary) warrant with respect to mathematical results involves some understanding of something to be said in favor of them or something about their further mathematical uses.

This is a familiar issue in debates between externalists and internalists about epistemic warrant. Although I endorse a qualified externalism as regards both interlocution and perception, I am doubtful about an unqualified externalism that would hold that if one had a reliable capacity to be smitten with answers to difficult mathematical problems, then one would be rationally entitled to rely on that capacity purely because of its *de facto* reliability. The case that I considered in the text is not quite so unqualified. The recipient can check the source's individual answers to difficult problems, to determine that each is correct. But it is hard for me to see how the recipient can be rationally entitled to rely on the special power of the source in the absence of an induction on these successes. I am inclined to think that the recipient's "knowing

how” to use the source is too thin, too much like the uncomprehending idiot savant’s power, to count as a rational entitlement (in this domain). I want to leave the issue open, however. It deserves deeper exploration.

29. These aspects, or notions, of objectivity are discussed in my “Individualism and Self-Knowledge” *The Journal of Philosophy* 85 (1988), esp. pp. 657ff.; and “Reason and the First-Person” forthcoming in *Self and Self-Knowledge*, MacDonald, Smith, and Wright eds.
30. George Boolos reminded me of Frege’s apt ridicule of Schroeder’s Axiom of Symbolic Stability, the idea that our justification in understanding a proof must invoke an empirical postulate about the physical stability of symbols used to express the proof. Cf. Gottlob Frege, *Foundations of Arithmetic* (Evanston, Illinois; Northwestern University Press, 1968), Introduction, p. viii. Frege was fully aware that use of one’s senses and the physical stability of symbols might well be psychologically necessary conditions for understanding mathematical reasoning over time.
31. A similar point would apply to our understanding of what is said. If reidentifying words required a complex empirical inference, then understanding would be interpretation, not comprehension. It would depend on reasoning *about* the words and their relation to what they expressed. They would be objects of reference in one’s understanding. In such cases, I think understanding itself might rest for its warrant on empirical reasoning rather than the direct application of an intellectual ability, given perceptual triggering. In such a case, one could not have an apriori entitlement to rely upon one’s linguistic comprehension of particular events with intentional content.