30

83690

UMI

MICROFILMED 2003

# Relying On Reason: A Reliabilist Account of A Priori Mathematical Knowledge

by

Mark V. McEvoy

A dissertation submitted to the Graduate Faculty in Philosophy in partial fulfillment of the requirements fo the degree of Doctor of Philosophy, The City University of New York.

2003

UMI Number: 3083690

UMI®



ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

This manuscript has been read and accepted by the Graduate Faculty in Philosophy in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

4/16/03
Date

Michael Devitt
Chair of Examining Committee

4/29/03
Date

Executive Officer

Jerrold J. Katz (Advisor, 1999-2002)

David M. Rosenthal (Acting Advisor, Fall 2002)

Michael Devitt

Michael Levin
Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

This manuscript has been read and accepted by the Graduate Faculty in Philosophy in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

# Abstract


Relying on Reason: A Reliabilist Account of A Priori Mathematical Knowledge

by

Mark V. McEvoy

Advisor: Professor Jerrold J. Katz

Acting Advisor, Fall 2002: Professor David M. Rosenthal

Because mathematical Platonism construes mathematical objects as existing outside of space and time, it precludes their having any causal interactions. This has led some to object that mathematical Platonism cannot explain how we know anything about such objects.

Process reliabilism sometimes evokes the converse objection. Since process reliabilism takes knowledge to be reliably produced true belief, it is sometimes said that the theory cannot explain the reliability of our mathematical beliefs, as we cannot interact causally with mathematical objects. If this were true, process reliabilism would be inadequate as a theory of knowledge, since it could not account for our mathematical knowledge.

My dissertation seeks to reconcile Platonism and process reliabilism by explaining the reliability of the processes which lead to our a priori knowledge of mathematics. Chapter One outlines and defends my conception of process reliabilism from Bonjour's counterexamples, from Feldman's "Generality Problem" and from evidentialist objections. Chapter two defends the notion of

the a priori from Quine's arguments. Chapter Three offers a conception of mathematical proof according to which it is a reliable, a priori process. It defends this conception of proof from attacks due to Kitcher and Tymoczko. Chapter Four argues that, contra Casullo, reliabilism can answer Benacerraf's challenge, and provide a satisfactory account of mathematical knowledge. Chapter Five offers a reliabilist account of the a priori process of mathematical intuition, and defends it against various objections.

Acknowledgments

I wish to acknowledge the enormous contribution of my advisor Jerry Katz, both to this dissertation, and to my approach to philosophy generally. I have benefited greatly both from his teaching, and from his written work. His encouragement, careful reading, considered criticisms and many helpful suggestions helped to make this dissertation a vastly better philosophical work than it otherwise would have been. His friendly concern for my progress was a valuable source of support to me while I worked on this project, and it was with much sadness that I learned of his death, early in 2002.

Special thanks are also due to David Rosenthal for taking over as my advisor in the Fall of 2003. David gave selflessly of his time, and put in an enormous amount of work in preparing me both for my defense, and for the job market. For this, and for the suggestions he made about the dissertation itself, I am hugely grateful.

This dissertation has also benefited greatly from the input of Michael Levin, whose speed in responding to drafts of individual chapters was matched by his microscopic attention to detail. Even the most casual glance through the footnotes of this work will reveal how much this dissertation owes to him. Additional thanks are due to Michael Devitt, who, although he didn't agree with a word of my dissertation, was nonetheless very generous with his comments, and to Phillip Kitcher, whose willingness to discuss his own work with me was very helpful at several stages along the road.

Thanks are also due to audiences at Columbia University Graduate Student Philosophy Conference, Kent State University Graduate Student Conference in Philosophy, CUNY Graduate Center Graduate Conference in Philosophy, CUNY Graduate Center Graduate Student Colloquium, CUNY Graduate Center Cognitive Science Colloquium Series, and the New Jersey Regional Philosophical Association.

I would also like to thank my family for their emotional, intellectual and financial support throughout my graduate studies. Particular thanks are due to my father, Tony McEvoy, my mother, Phil McEvoy, my brother, Keith McEvoy, my niece, Jade McEvoy, and my fiancée, Stephanie Sapiie. I would also like to thank my friends in New York, both in the Philosophy Program, and outside of it, for their support during my time here.

Table of Contents

# Introduction

This dissertation defends a process reliabilist account of mathematical apriorism, on the assumption of Platonism. It is thus, on the one hand, an apriorist solution to the objection that Platonism is epistemologically bankrupt, and, on the other, a defense of reliabilism, both against the objection that it is incapable of accounting for mathematical knowledge on the assumption of Platonism, and against the many other objections aimed at that theory.

The seeds of this dissertation were sewn when I first came to the CUNY Graduate Center in the academic year of 1996-7. In that first year there, I took classes with Jerrold J. Katz—who sadly passed away as I was finishing this project—and Michael Levin. These classes, I now see, started me down the road of writing this dissertation. Katz's stalwart defence of such philosophically unfashionable concepts as analyticity, apriorism, and Platonism intrigued me. Moreover, his level-headed approach to apriorism was instrumental in my coming to believe, as he did, that many of the objections to apriorism could be met by providing a suitably demystified conception of the a priori. Later, when I came to read his Realistic Rationalism,[1] I saw in the mathematical apriorism defended there hints of a reliabilist flavor.

Levin's class on reliabilism provided an excellent exegesis, and spirited defense, of that school of epistemology, thereby developing my own nascent interest in it. It was

[1] Jerrold Katz, Realistic Rationalism (Cambridge, Mass.: Bradford MIT, 1998).

in this class that I first encountered Benacerraf's paper "Mathematical Truth,"[2] and the challenge there mounted to mathematical epistemology. The problem, simply put, is that it appears impossible to hold both our best theory of mathematics—Platonism—together with a plausible epistemology. The problem has two aspects. The problem for the Platonist is to show how knowledge of abstract objects is possible, in the absence of any causal connection between us and those objects. The problem for the reliabilist, as developed by Albert Casullo,[3] is that he can offer no account of our mathematical knowledge, when this knowledge is construed as knowledge of abstract objects, wholly lacking in causal powers.

In a sense this dissertation is an extended answer to the Benacerraf problem. Since, however, it also offers extended defenses of reliabilism, a priori knowledge, and mathematical apriorism, it is more than just this.

**A Quick Sketch of the Main Characters**

**a) Reliabilism**

Reliabilism, as I shall understand it, is a theory of knowledge, which holds that knowledge is true belief, produced or sustained by a reliable process. Although usually conceived of as an empiricist theory of knowledge, since it does not require of a reliable process that it be one which causally links the object known with the knowing subject, reliabilism is, on the face of things at least, quite compatible with apriorism.

---

[2] Paul Benacerraf, "Mathematical Truth, in Philosophy of Mathematics, 2nd ed. , ed. Paul Benacerraf and Hilary Putnam (Cambridge University Press: Cambridge, 2cd. Edition, 1983).

[3] See his "Causality, Reliabilism, and Mathematical Knowledge," Philosophy and Phenomenological Research, September (1992): 557-84.

Reliabilism pointedly does not require that a subject's warrant for a given belief be consciously accessible. It thus allows for the possibility that a subject might know some p without knowing that he knows it. For the reliabilist, once the producing (or sustaining) process is both reliable and repeatable, the belief is an instance of knowledge.

**b) Platonism**

Platonism, as we shall be concerned with it here, is a view of the ontology of mathematics, according to which mathematical objects exist abstractly and independently of minds. These objects are the truth-makers of mathematical propositions, and the question of whether a given mathematical proposition is true is settled by whether it correctly represents mathematical reality. The form of Platonism I shall be discussing takes mathematical objects to be aspatial, atemporal and acausal. It also takes mathematical truths to be necessary.

**c) Mathematical Apriorism**

Mathematical Apriorism is an epistemological thesis, which holds that the truths of mathematics are knowable a priori. A priori knowledge is knowledge the warrant attaching to which does not depend upon any particular characteristic of experience. Although the proofs of logic, to take one example, typically are perceived, the warrant attaching to the conclusion of such a proof depends not on the fact that the proof was seen, heard, or otherwise perceived, but on the fact that the conclusion of the proof must be true, granted the validity of the logical form together with the truth of the premisses. Thus, in the typical case, I would have a priori warrant for my belief in the conclusion of a given proof once I understand that proof.[4] Of course, I usually perceive the proof prior to understanding it, but this perception is not itself the warrant, it is only the access to it.[5] The process of following a mathematical proof, then, will properly be counted as a priori if it does not require experience for its warrant—provided of course, that we can defuse those reasons offered for denying that mathematical knowledge is a priori knowledge. If mathematical processes are also repeatable and reliable, they will, by reliabilist lights, count as a priori warranting process.

This is all somewhat vague, and what it means for a mathematical process to be a priori in this sense will, of course, be sharpened over the course of this dissertation, but what we have said so far is notable for what it does <u>not</u> require of the a priori. According to the conception of the a priori to be defended here, putative items of a priori knowledge—that is, beliefs produced by a priori methods—are neither infallible nor immune to rational revision. For a reliabilist, an a priori method will be warrant-conferring if it produces true beliefs most of the time. One hundred per cent reliability is not required.

A priori methods are rationally revisable. If other a priori reasons are advanced which show a putative a priori truth to be mistaken, one ought, rationally to revise one's

---

[4] I say "in the typical case" here because I will be discussing cases of a priori knowledge in which the subject has no access to his warrant, and cases in which the proof of a given theorem exists only in a computer.

[5] This account of the a priori, and the distinction between warrant and access to that warrant is taken from Paul Boghossian and Christopher Peacocke's introductory essay to their <u>New Essays on the A Priori</u> (New York: Oxford University Press, 2000).

belief. They are, however, not subject to revision for empirical reasons. An a priori truth, then, is one which is knowable via an a priori method. It is thus possible, on this understanding of the a priori, both for an a priori truth to be known a posteriori, and for a truth to be an a priori one, even if nobody were actually to know the truth in question. Once it would be possible for someone to come to know the truth in a way which did not depend, for its warrant, upon experience, I count this truth as a priori. This point will loom large in my Chapter Three discussion of Tymoczko.

**Making The Platonist's Problem Disappear: Empirical Knowledge of Abstract Objects**

One way round the Platonist's problem is to deny the impossibility of empirical knowledge of Platonistic objects. One might, with an earlier time-slice of Penelope Maddy,[6] deny the claim that mathematical objects have no causal powers. Or, one might instead, with W. V. O. Quine[7] and W.D. Hart,[8] accept that such objects have no causal powers, but deny that this prevents us from having (a posteriori) inferential knowledge of such objects, since our best science must posit such objects. We have knowledge of abstract objects on this picture, since they are indispensable to our best science, and since our best science is known to be true on the basis of its success in describing,

---

[6] See her Realism in Mathematics (Oxford: Oxford University Press, 1990).

[7] See, for example, Quine's "Two Dogmas of Empiricism," reprinted in his From a Logical Point of View, 2nd. ed., (Cambridge, Mass: Harvard University Press, 1961).

[8] W.D. Hart, "Access and Inference," in The Philosophy of Mathematics, ed. W.D. Hart (New York, Oxford University Press, 1996), 60 – 61.

predicting, and controlling nature. Both of these solutions are consistent with reliabilism, but neither with a priori reliabilism.

These empiricist brands of Platonism, however, are open to serious—possibly insurmountable—objections. Maddy's version of the theory (she has since rescinded it) holds that we do interact with mathematical objects. We interact with a six-membered set, for example, when we purchase a half-dozen eggs. Sets and numbers, that is, are concrete objects, and are capable of having physical location. The claim that one actually sees the number six when one opens the fridge is a highly counter-intuitive one. But falling the wrong side of intuition is not the worst of it for this position. If mathematical objects are concrete, how can the infinitude of mathematical objects be accounted for? Where, on this position, is the null set? If it is not anywhere, how can it be concrete; if it is not concrete, what is it?[9] The position, as Maddy herself has recognized, is unpromising

Nor does Quine's brand of empiricist Platonism fare much better. As we shall see in Chapter Two, such a position cannot account for the truth values of unapplied portions of mathematics. It also has the counterintuitive consequence that research in mathematics must wait upon the physical sciences for its warrant. Finally, there is the objection, raised by several authors,[10] that Quine's epistemological holism is incoherent.

---

[9] For more on these and other objections, see Katz (1998), and Charles Chihara, "A Gödelian Thesis Regarding Mathematical Objects: Do They Exist? And Can We Perceive Them?" The Philosophical Review, 91 (1982): 211-27.

[10] See, for example, Katz (1998), 72-74 and Richard Warner, "Why Is Logic A Priori?" The Monist 72 (1989): 40-51.

**Making the Reliabilist's Problem Disappear: Rejecting Platonism**

The reliabilist might attempt to avoid the problem of accounting for mathematical knowledge by simply denying the truth of Platonism. This move is less than inviting, for at least two reasons. Since one of the objections to reliabilism is that it cannot account for our knowledge of mathematics, when the objects of mathematics are construed Platonistically, it would be far better for the reliabilist to show how knowledge of such objects is possible, rather than simply denying Platonism. This would provide an answer to the objection, instead of initiating a degenerating dispute where each party simply denies the other's premiss. More importantly, Platonism has strong credentials as a mathematical ontology, and it would be wise for the reliabilist not to tie his fate to mathematical nominalism. Far better for him to be able to say how we come to have mathematical knowledge whatever the correct ontology of mathematics might be.

**Addressing the Problems: Combining Platonism, Apriorism and Reliabilism**

I shall not, then, be rejecting Platonism. However, I shall be assuming, rather than arguing for, the truth of Platonism throughout this dissertation. This is for two reasons. From the reliabilist perspective, the problem is to reconcile his account of knowledge with a Platonist ontology. Without the assumption of Platonism, the problem disappears. From the Platonist's perspective, the problem is to show how knowledge of abstract mathematical objects is possible. Again, the problem only arises on the assumption of Platonism. In short, I assume Platonism because, without this assumption, there are no problems to solve.

Since I will be assuming Platonism, and since I have already said that one aspect of this Platonism is that mathematical truths are necessary, this assumption of Platonism involves an assumption of the necessity of mathematical truths. This assumption of necessity will play a role in Chapter Three's discussion of Kitcher's anti-apriorist arguments,[11] but it will not be argued for here. There are two reasons for this. The first is that since the problems I shall be discussing arise only on the assumption of Platonism, and since Platonism traditionally holds mathematical truths to be necessary, to assume Platonism simply is to assume that mathematical truths are necessary. The Quinean brand of contingent Platonism, which arises out of his epistemology of empirical holism, I believe to be open to serious objections. I shall discuss these objections in Chapter Two, and in doing so will, in effect, present reasons for thinking that the only tenable brand of Platonism requires necessity.[12]

The second reason is that Kitcher's arguments against apriorism are meant to show that there is no defensible brand of mathematical apriorism. The apriorist is thus permitted to show how Kitcher's arguments fail against a Platonist conception of mathematical apriorism—one which includes necessity.

---

[11] These arguments are to be found in Phillip Kitcher's, The Nature of Mathematical Knowledge (New York, Oxford, 1983), and in his "A Priori Knowledge Revisited," in Boghossian and Peacocke (2000): 65-91.

[12] I shall not be addressing Quine's arguments that there are no necessary truths. I have quite enough to do here as it is, without entering that debate. For responses to Quine's arguments, see Ruth Marcus's "A Backwards Look at Quine's Animadversions on Necessity," in Perspectives On Quine, ed. Robert Barrett, and Roger Gibson (Oxford: Basil Blackwell, 1990): 230-243. For some more general considerations in favour of the existence of necessity, see Saul Kripke's Naming and Necessity, Cambridge: Harvard University Press, 1980).

With the assumption of Platonism, the reliabilist is left with the following situation. The attempt to answer the Platonist's epistemological problem by adopting mathematical empiricism runs into its own problems. The reliabilist who takes realism seriously, and who wants to answer the objection that his theory cannot account for mathematical knowledge, would do well to develop an account of such knowledge according to which it is a priori.

Given a viable account of the a priori, it might appear that there is a very straightforward reliabilist solution to the problem of mathematical knowledge. For the process reliabilist, knowledge is simply true belief, caused by a reliable process. Thus, given that we have mathematical knowledge, there must be a reliable process which yields that knowledge, but which does not require the impossible causal link between ourselves and a platonic realm. Reliabilism, that is, since it requires only a reliable process, and not one which involves a causal link between the knowing subject and the object of knowledge, seems completely amenable to an a priori solution.

However, as we shall see in Chapter Four, Albert Casullo (1992) has argued that this advantage of reliabilism is merely apparent, and that no tenable version of that theory can account for knowledge of abstract objects. Moreover, the "solution" is too easy. If one says that there is some reliable process which yields knowledge of a mathematical realm, then, even if one adds that this process is not a causal one, one merely invites the challenge to describe the process.

The problem, then, is to combine a priori Platonism with a reliabilist epistemology, in such a way as to answer the challenge faced by each theory. For the

Platonist, that challenge is to produce a viable epistemology for his ontology. For the reliabilist's part, he must answer the objection that he cannot account for mathematical knowledge on what, to many philosophers, appears to be the strongest mathematical ontology. As a reliabilist who thinks there is a priori knowledge, and who is usually persuaded by Platonism, I would like to be able to consistently hold Platonism, apriorism and reliabilism. To my knowledge, however, there is no position which attempts to solve the problems faced by each theory by combining those theories.[13] I intend to remedy this deficiency.

**More Problems**

Prior to answering the problems arising from the assumption of Platonism, however, some other concerns must be addressed. Many philosophers find one or other, or both, of reliabilism and apriorism to be extremely implausible. Important objections to both are to be found in the literature. The situation with these aspects of my position is unlike that with Platonism. I cannot decline to defend these aspects of my position by observing that, without assuming their truth, the problems do not arise. Even the empiricist Platonist must address the objection that his position has no satisfactory mathematical epistemology, and it is not only the reliabilist who must face up to the objection that his epistemology makes it difficult to see how we could have knowledge of abstract objects. Since, then, the problems I am addressing arise independently of reliabilism and apriorism, challenges to these theories must be addressed, lest the

[13] While Katz's work has a reliabilist flavor, he remained an internalist about knowledge, and thus was not a reliabilist. While he did remark to me that he had considered calling his mathematical epistemology a reliabilist one, he also stopped short of doing this.

attempt to provide a satisfactory mathematical epistemology by combining these theories be deemed futile on other grounds.

The order of business will be as follows. In Chapter One, I articulate my conception of process reliabilism, according to which, knowledge is true belief, produced by a reliable process, and defend it against various objections. In Chapter Two, I outline the understanding of the a priori which I shall use here, and defend it against Quine's attacks on the a priori

The remaining chapters defend the traditional conception of mathematical knowledge as knowledge a priori, from within a reliabilist framework. There are two cases: basic—or non-inferential—mathematical knowledge, and inferential mathematical knowledge. One might think that if one could present a compelling defense of the claim that our basic mathematical knowledge is a priori, the second case would be rather easy. Given that our basic knowledge is a priori, then, since the rest of our mathematical knowledge is deduced from this via rules of inference which are themselves either part of our basic mathematical knowledge, or else are logically warranted, it would seem to follow that the rest of our mathematical knowledge must be a priori. At least, it would seem to follow if one takes logic to be a priori. Since, by this time, I will have offered reasons to reject Quine's claim that all knowledge is a posteriori, the main reason for holding that logic is not a priori will have been removed. In any event, it is commonly thought that the crux of the matter for the apriorist is his account of basic mathematical knowledge.[14] However, there are objections to the

[14] Hartry Field, for example, remarks that the Platonist can make use of the axiomatization of mathematics to account for his alleged knowledge of non-basic mathematical propositions, and that the crux of the issue is then his alleged knowledge of basic mathematical propositions. See his "Realism,

apriority of mathematical proof due to Philip Kitcher (1983 and 2000) and Thomas Tymoczko[15] which focus on very long proofs. Kitcher's argument is based on the limitations of human memory, while Tymoczko's is based on the use of computer proofs in mathematics. In Chapter Three, I address these arguments, along with some additional arguments presented by Kitcher (2000), and conclude that, according to the moderate conception of the a priori briefly sketched above, mathematical proof is an a priori method, which yields a priori knowledge.

In Chapters Four and Five, I articulate and defend a reliabilist account of mathematical intuition. Chapter Four addresses Albert Casullo's argument that no credible form of reliabilism will be able to offer an account of a priori mathematical intuition which can solve the Benacerraf problem. In Chapter Five, I offer a Katzian account of a priori mathematical intuition, within a reliabilist framework, and show how it withstands arguments of Kitcher's to the effect that no process of intuition could produce a priori warrant.

Mathematics, and Modality," in Realism, Mathematics, and Modality (Oxford: Basil Blackwell, 1989): 231.

[15] Thomas Tymoczko, "The Four-Color Problem and its Philosophical Significance," Journal of Philosophy, Vol. LXXVI, No.2 (February 1979): 57-82.

# Chapter One

# Reliabilism

## Introduction

This chapter offers an outline of what I take to be the most plausible form of reliabilist epistemology, and goes on to defend that theory from several important objections, thereby rendering our initial outline more detailed. The version of reliabilism which emerges from this chapter will make up the epistemological framework for the rest of the dissertation. And, in later chapters, it will be from within this framework that the problems for mathematical apriorism, and for Platonist epistemology, will be addressed.

## 1.1

## A Brief Outline of Reliabilist Epistemology

Reliabilism is an externalist theory of knowledge with two main features. The first of these is the analysis of knowledge it offers. According to the reliabilist, one knows that p when the following conditions are satisfied:

(1) p is true

(2) One believes that p, and

(3) One's belief that p is caused or sustained by a reliable belief forming process.

The second main feature of reliabilism, a feature shared by all externalist accounts of knowledge, is a consequence of this analysis. This is a rejection of the KK thesis (the thesis which states that a necessary condition for knowing any proposition is that one knows that one knows that proposition).

A consequence of these two features of reliabilism is that reliabilism neither is nor is supposed to be criterial for knowledge. There is, in general, no reason to expect that one can tell, from the inside, as it were, that any given belief which one might have actually satisfies the conditions for knowledge. This, intuitively, is simply what it is for a theory of knowledge to be characterizable as externalist. Unless one's beliefs are "hooked up" to the world in the right way, one's beliefs will not be instances of knowledge. Nor, according to the externalist, is there any compelling reason to think that one will necessarily be aware that one's beliefs in fact are so hooked up.

The hapless envatted brain provides an extreme example of both points here. Firstly, the brain's "experiences" are supposed to be indistinguishable from those genuine experiences of regular folk. But the brain's beliefs, including any true beliefs the vat operators cared to induce, do not constitute knowledge. It seems clear that the reason for this is that those beliefs are not hooked up to the world in the right way.

Secondly, the epistemological puzzle raised by envatted brains is that of explaining how we could tell that we are not similarly deluded. Since this seems to be an exceptionally difficult puzzle to solve, it gives credibility to the externalist's claim that there is no reason to think that we will necessarily be aware that our beliefs are correctly hooked up to the world. Which is to say that there is no reason to expect an analysis of knowledge to be criterial for knowledge.

Whether the reliabilist analysis of knowledge is satisfactory depends in large part on how the key terms in (3) are fleshed out. In particular, one wants an account of what is meant by the term 'reliable process'. The remainder of this chapter will be devoted to describing such an account, to outlining the motivations for reliabilism, and to dealing with the most common objections to the theory.

## 1.2

## Why Would Anyone Be a Reliabilist?

### 1.2.1.

### Post-Gettier Epistemology and the Failure of the CTK

In the immediate aftermath of Edmund Gettier's paper, "Is Justified True Belief Knowledge?"[16] the causal theory of Knowledge (CTK) gained credence. It did so in part because of its ability to deal with Gettier cases[17] and partly because of the obvious causal element in certain knowledge-yielding processes (e.g., perception). CTK, however, fell rather quickly into disrepute. There were problems with spelling out what was meant by an "appropriate" causal connection between the fact known and the knower (if a benevolent demon caused my belief in an actually existent external world,

---

[16] Edmund Gettier, "Is Justified True Belief Knowledge?" Analysis, Vol. 23 (1963): 121-3.

[17] The thrust of the CTK's analysis of Gettier cases is that there is no causal connection between (e.g.) the fact that someone in the office owns a Ford, and Smith's belief that someone in the office owns a ford. See Alvin Goldman's "A Causal Theory of Knowing," Journal of Philosophy 64 (1967): 357-72.

does that belief count as knowledge?). Counterexamples to the theory proliferated.[18] Moreover, in order to account for mathematical and moral knowledge, CTK had to take a controversial stand on realism, since neither mathematical objects nor moral truths are possessed of causal powers.[19]

These problems paved the way for reliabilism. Reliabilism, it seemed, could save what was good about CTK without succumbing to its problems. Reliabilism could account for perceptual knowledge since perception is, generally, a reliable method of forming true beliefs. From a reliabilist perspective, what makes perceptual knowledge knowledge, is not the mere fact that it was caused in an appropriate way. It is, rather, that the causal process which underlies perceptual belief is reliable—it reliably leads to true beliefs. Put another way, what made CTK seem like a good theory in the first place, was the reliability of the causal processes in the examples used to motivate CTK.

Reliabilism can account for Gettier cases just as well as could CTK. Jones, on the basis of seeing Brown in a Ford forms the belief 'Either Brown owns a Ford or Jones is in Barcelona'. As it happens, Brown is in a borrowed Ford, and, unbeknownst to Jones, Smith has coincidentally just arrived in Barcelona on vacation. Although Jones's disjunctive belief is true, it is so merely fortuitously, since Jones has no reason at all to believe anything about Smith's whereabouts. The disjunctive belief, then, is not

[18] To cite just one, there is the case of Henry and the barn facades. Henry is in an area where there is a preponderance of fake barns. Henry sees an actual barn, and forms the belief that he is looking at a barn. Since, under the circumstances, Henry might well have been looking at a barn facade, we are reluctant to attribute knowledge to Henry. However, since Henry's belief is caused by his seeing an actual barn, the CTK yields the wrong result here. See Alvin Goldman, "Discrimination and Perceptual Knowledge," in <u>The Theory of Knowledge</u>, ed. Pojman (CA.: Wadsworth, 1999): 169-182.

[19] See Mark Steiner's "Platonism and the Causal Theory of Knowledge," <u>Journal of Philosophy</u> 70 (1973): 57-66, for an attempt to defend the CTK against this objection.

knowledge. According to the reliabilist this is due to the fact that inferring a disjunction from a falsehood is not a reliable process.

Moreover, the causal inertness of mathematical objects and moral truths presents no immediate problem for reliabilism. Since reliabilism does not require a causal link between the fact known and the knower, the problem which beset CTK does not arise here. On a reliabilist picture, if our true mathematical and moral beliefs are arrived at via a reliable process, then they are instances of knowledge, regardless of whether the object of knowledge is or is not causally inert.[20]

The immediate post-Gettier attraction of reliabilist epistemology, then, is that it can account for Gettier cases without running into the problems which led to the decline of CTK.

### 1.2.2

### The Attractiveness of the K/KK distinction

Consider the reliabilist conditions (1) – (3) above. Take any p which does not include epistemic terms—say 'That is a tree'—for which (1) - (3) are satisfied. S will count as knowing that p, according to the reliabilist. However, because condition (3) requires only that the belief be produced by a process which is in fact reliable, and not that S know (or believe) that the process is reliable, S will not thereby count as knowing that he knows that p. S will count as knowing that he knows that p only if (1) - (3) are satisfied when '<u>I know</u> that that is a tree' is taken as the value for p. As it is not

[20] The question of whether this apparent advantage of reliabilism survives close scrutiny is the subject of Chapter Four.

generally the case that one knows that one's belief is produced by a reliable process, one can know that p without knowing that one knows that p.

Reliabilists insist on this distinction for several reasons. Firstly, it is clear that the propositions 'p' and 'I know that p' are different propositions. At a minimum, they have different truth-conditions. Since they are different propositions, it is surely possible to believe one and not the other. One way this might happen is as follows. As the two beliefs are different, it is conceivable that the beliefs be produced by different processes. Granted this possibility it is then possible that the processes be independent of each other, and one might be led, by a reliable process to form the belief that p, without there being any operative process leading to a belief that one knows that p.[21]

A second reason why reliabilists insist on the distinction is that the requirement that one must know that one knows in order to know is too restrictive, and would rule out much that we would want, intuitively, to call knowledge. Small children, and animals for example, surely know things. It is at best improbable, however, that they have higher-order beliefs which they could use to justify their knowledge. They would not therefore satisfy the condition that they know that they know, and so, if this condition is criterial for knowledge, they do not count as knowing.

Nor would it be much help to attempt to blunt the force of this point by holding that neither animals nor small children know things in the same way that we do.[22] Clearly there are differences between the way they know things, and the way we do, but

---

[21] Consideration of Moore's paradox suggests that one could not both believe that p and believe that 'I know that p' is false, where the value of p is constant. This is quite consistent with my point in the text.

[22] Thanks to my friend Matt Kaiser for raising this issue with me.

we would still require an account of what their knowledge has in common with our knowledge, and which legitimizes ascriptions of knowledge (as opposed to belief) to such beings. Absent such an account, the KK requirement places these beings outside the class of knowers.

A further, though related, point is that the requirement that one know that one knows rules out more than just sub-linguistic creatures from the class of knowers. There are many cases in the contemporary literature where one intuitively wants to say that the subject knows a given fact, but the KK requirement does not allow such a verdict. Take Nozick's telephone case, for example. Knowing that his patient would be devastated if his forgetful wife does not call him to wish him a happy birthday, the analyst in the case has hired a standby actress to call the patient and pretend to be his wife, just in case his wife forgets to call him. As it happens, the patient's wife does call him.[23] Intuitions are divided here, but it seems to me that we want to say that the patient knows he is talking to his wife, even though he would have been taken in by the actress. The reliabilist can account for this intuition. The patient's belief is true, and it is formed by a reliable belief-forming process (assuming that it is not a regular occurrence that there be a paid actress on standby every time his wife is supposed to call). He therefore knows that he is talking to her. On the other hand, he might well not be able to tell that the process is reliable, or to give an account of how it is that he knows that it is his wife to whom he is speaking. He does not, therefore, know that he knows.

---

[23] Robert Nozick, Philosophical Explanations (Cambridge, Massachusetts: Belknap, Harvard University Press, 1981): 190. Nozick's own analysis yields the (in my view, mistaken) conclusion that the patient does not know that he is talking to his wife, since he would believe he was even if he were not.

Another such case is due to Goldman, and here the verdict seems more clear cut. The KK thesis requires that a knower be able to cite his justification for his knowledge. But for many beliefs which seem, intuitively, to be instances of knowledge, we have forgotten the justification for these beliefs. Goldman's example is of his coming to believe that Lincoln was born in 1809 through reading it in an encyclopedia, but later coming to forget having read of this. It seems clear that Goldman knows when Lincoln was born, even though he cannot cite his justification. The criterion imposed by the KK thesis, however, is not satisfied.

A further problem with the distinction is that it seems to involve us in an infinite regress. If, in order to know that p, I must first know that q, where q is the proposition that I know that I know that p, how can I come to know that q? Clearly, I must first know that I know that q; that is, prior to simply knowing that p, I must first know that I know that I know that p. In a nutshell, the problem is that since any proposition can be so embedded, to require that one know the embedding proposition before one counts as knowing the embedded proposition will mean that the requirements for knowledge will always be out of reach. Reliabilism, by rejecting the KK requirement, cuts off this regress at the first juncture.

Perhaps the most compelling reason for insisting on the distinction, however, is that, given this distinction, the externalist epistemologist is in a position to respond to one popular type of skeptical argument. The argument in question runs as follows. In order for you to know (e.g.) that there is a book on that desk, you must first know that the entire scene is not some elaborate hoax dreamt up by a Cartesian demon. Since you do not know this, you do not know that there is a book on that desk. A similar argument shows that you do not know any of the things you take yourself to know, and global skepticism results.

However, as Alston notes,[24] while it is true that one would not know that there was a book on the desk if the whole scene were an elaborate illusion, this does not show that one does not know that there is a book on the desk if the whole scene is not an illusion. This latter follows only on the assumption that, in order to know that there is a book on the desk, one must first know that the scene is not an illusion. But to require this assumption is to require that one know that one knows that there is a book on the desk, and it is unclear why one must know this fact in order to know that there is a book on the desk. For sure, the possibility exists that my current perceptual belief is the result of some abnormal process. But, as Alston asks, if in fact my belief is caused normally, why should the fact that I do not know this prevent me from knowing that there is a book on the table? More generally, why must I be required to know something which, epistemologically, is one level up, in order to know something on an epistemologically lower level? The skeptic, from a reliabilist perspective, errs in requiring that we know that we know before allowing that we know.

The example of global skepticism serves also to cast doubt upon the credentials of the internalist demand that the KK thesis be satisfied. The victim of the Cartesian demon takes himself to know certain things about the (simulated) external world, just as we take ourselves to know various things about our world. If internalism about knowledge were correct, the Cartesian dupe ought, in principle, to be able to tell that he

[24] William Alston, "Level Confusion in Epistemology," Midwest Studies in Philosophy V (1980): 135-150.

does not know anything about the world. It would be to put things mildly to say that it seems difficult to imagine how he could come to discover this.

Considerations such as these show that the KK condition is overly restrictive. Reliabilists therefore insist on a distinction between knowing and knowing that one knows.

## 1.3

## Objections to Reliabilism

This section does double duty. Firstly, it is an attempt to provide answers to the most important objections to reliabilism. Secondly, in so doing, it will bring into sharper focus the rather vague picture of reliabilism offered in the opening section. The objections to be considered here are, in order of appearance, Bonjour's counterexamples, the generality problem, and finally, the evidentialist claim that the reliabilist account of warrant is either mistaken or redundant, depending on how it is spelled out.

### 1.3.1

### Bonjour's Counterexamples to Reliabilism

In The Structure of Empirical Knowledge,[25] Lawrence Bonjour presents a number of alleged counterexamples to externalist epistemologies. I will deal here with two well known (and representative) cases, the case of Maud, and that of Norman. Maud thinks

[25] Lawrence Bonjour, The Structure of Empirical Knowledge (Cambridge, Massachusetts: Harvard University Press, 1985).

that she is clairvoyant, though she has no reasons for this belief. She perseveres with this belief despite being in possession of quite compelling scientific evidence to the effect that clairvoyance is impossible. One day, Maud comes to believe that the President is in New York, though she has no evidence to support this belief. She claims that the belief is a result of her psychic powers. It turns out that the President is in New York, and that Maud does have completely reliable psychic powers, and furthermore that her belief about the President was produced by these powers.

Bonjour concludes that the reliabilist is committed to saying that Maud's belief is justified, whereas in fact, Maud is not justified in her belief, because she has excellent reasons to believe that clairvoyance is not possible. For the reliabilist who wishes to avoid such cases by explicitly including a further condition on justification and knowledge—to the effect that S must not have good evidence that the process which has produced his belief does not exist (as will emerge, I am not such a theorist)—Bonjour presents the case of Norman.

Norman, like Maud, is clairvoyant. Unlike Maud, he possesses no evidence for or against the possibility of such a phenomenon. Nor does he have any evidence for or against the thesis that he possesses it. He forms the belief that the President is in New York by means of his clairvoyance, and the belief is in fact true. Again, according to Bonjour, Norman's belief does not count as knowledge.

A preliminary question which ought to be asked is whether the alleged counterexamples are coherent. Can Maud and Norman actually have the beliefs ascribed to them by Bonjour? Take Maud's alleged belief. She is stipulated to be in possession of compelling scientific evidence to the effect that clairvoyance is impossible. It is at least problematic to claim that she has all this evidence, and none at all for her belief in her own clairvoyance, but yet still believes that she is clairvoyant. To simply accept this counterexample at face value is to assume that it is possible for someone's belief that p to be maintained in the face of overwhelming evidence—of which the subject is aware—to believe not-p, and no reason at all to believe that p. This assumption is at least questionable, and sufficiently so as to cast doubt on the view that the Maud case is devastating to reliabilism.

Norman is even worse off. Unlike Maud, he doesn't even believe that he is clairvoyant. His "belief" just appears, unexplained and uninvited, from nowhere. Isn't it at least problematic to claim that one could maintain belief in some contingent, empirical, proposition, p, with no evidence at all for p, nor any idea at all as to how or why one has come to believe that p? Does it not seem more likely that once Norman became aware of the precariousness of his belief state, he would discard his belief as to the whereabouts of the President? It certainly seems so. And if it is so, it is again questionable how damaging the example is for reliabilism.

However, even if it is insisted that Norman and Maud can, and do, have the relevant beliefs, this does nothing to establish Bonjour's anti-reliabilist conclusions. So, for the sake of argument, it will be granted that they do have these beliefs.[26]

Having granted this, the question then arises as to whether, if Maud and Norman can believe, in the face of contrary evidence, the reliabilist is mistaken to claim that those beliefs are justified. The reliabilist can respond that a belief is justified if it is the

[26]The question of whether the belief is possible will become important when dealing with a complication added by Bonjour to the Norman case. However, independent argument will be produced at that point to urge the conclusion that Norman does not believe.

result of a reliable belief-forming process. Ex hypothesi, Maud's belief is the result of such a process, therefore her belief is justified. We are reluctant to say that her belief is justified, because, in our world, clairvoyance is not a reliable belief-forming process, and as a result, our intuition is to reject any attempt at justification which appeals to clairvoyance. However, it is stipulated in the case that clairvoyance is reliable, so Maud's belief is justified. She could not give a justification of the clairvoyant process, but that would be relevant to justifying her belief that she is justified in believing that the President is in New York. It is not relevant to her being justified in believing the President is in New York.

Compare Maud's case with that of idiot savants. Somehow, the idiot savant knows the answer to a great many questions of a certain sort. He does not, however, have access to reasons which would justify him giving the answers he does give. But we do not say, on those grounds, that he does not know. Compare also the case of Ramanujan. Hardy describes him as having, at best, a rudimentary grasp of mathematical proof.[27] However, he was able to grasp, intuitively, a wide range of complex mathematical truths. It seems churlish to deny knowledge of these truths to Ramanujan merely because he would have been unable to produce a complete formal proof with which to justify his answer to a given question. Similarly, given how the case is described, for Maud, she is clairvoyant. Her belief as to the whereabouts of the President is a result of this reliable clairvoyancy. The intuition which Bonjour attempts to fuel relies on us not playing by the rules of the case as described. It relies on us importing our real-world knowledge that beliefs formed by an alleged clairvoyant

[27] G.H. Hardy, A Mathematician's Apology (London: Cambridge University Press, 1940).

faculty tend towards truth about as reliably as those formed via wishful thinking. But to import our real-world knowledge that there is no such thing as clairvoyance is simply to disregard the rules of the case. The ground rules of both cases state that the protagonists are reliably clairvoyant. Why deny them knowledge simply because they lack access to the mechanism of their clairvoyancy?

Consider the following, structurally parallel, case. Suppose that Blinky lives in a world where scientists claim to have shown conclusively that there is no such sensory modality as vision. There is a large literature (presumably in Braille) to the effect that there is no serious evidence whatever that such a modality exists. Blinky is in possession of this evidence, but continues to believe himself inexplicably blessed with this mysterious fifth sense. In defiance of scientific consensus, he regularly forms beliefs by means of this modality. The modality is reliable, and yields true beliefs in an overwhelming number of cases. Suppose next that Blinky sees a book on his table (in good lighting conditions, from fairly close up, with no tricks), and forms a belief that there is a book on the table. Does his belief constitute knowledge? It seems obvious that it does. And this despite the claims of the scientists of his world.

The following parallels clearly hold between Blinky and Maud: Firstly, all reputable scientists hold that the faculty in question does not exist. Secondly, in each case, the gifted person is aware of this evidence, but weights their own personal experience more heavily. Finally, in both cases, the relevant modality is both reliable, and experienced as such, by the gifted person. If, as seems clearly true, Blinky's reliable faculty of vision regularly yields knowledge, on what basis can we hold that Maud's reliable clairvoyance does not?

The Blinky case, I conclude, shows two things. Firstly, it shows that the intuitions yielded by the Bonjour cases depend on us not playing by the rules of the case. Those cases thereby assume, rather than show, what they need to show in order to trouble the reliabilist. Secondly, the Blinky case shows that for a basic-belief-forming process to yield warranted beliefs, it is not required that the subject hold appropriate second-order beliefs about that process. It is enough that the process actually be reliable.

Bonjour goes on to claim that in order for a person's belief to be justified, the "believer in question must know or at least justifiably believe some ... set of premisses or reasons [in support of his belief], and thus be himself in a position to offer the corresponding justification"(Bonjour, 1985, 43). This, of course, is exactly what the externalist denies. Bonjour abuses externalism as simply waiving the above general requirement for justification in a certain class of cases. But it is not as if there is no principled reason for waiving the requirement. It is not that the reliabilist waives the requirement when faced with a case that he cannot otherwise deal with. The point is that <u>access</u> to one's justification is at a higher level, epistemically speaking, than justification <u>simpliciter</u>. And this is held universally true by the reliabilist. One may reject the externalist's account, as Bonjour does, but independent argument is needed if thus rejecting it is to be principled. Bonjour provides no such argument. In the absence of any such argument, and in light of the previously discussed advantages of the externalist position on the KK thesis, the reliabilist can, with theoretical justification, claim that Norman's belief is justified.

But, regardless of whether the beliefs themselves are justified, are the believers justified in holding them? That is, ought they to believe? It seems clear that neither Maud nor Norman ought to believe that the President is in New York, for they have no epistemic access to their justification. This conclusion need not embarrass the reliabilist, however. For, on the one hand, as we have seen, he can reject the claim that they can have the relevant beliefs. On the other hand, he can agree that they ought not to believe, but in fact do. If Maud and Norman were being epistemically responsible, then, given the respective evidence they have for and against their beliefs, they would give up those beliefs. As it happens, they do not. They cling to their beliefs with respect to the whereabouts of the President. Although they ought not to hold these beliefs, since those beliefs are, ex hypothesi, reliably produced, they constitute knowledge.

The jarring note produced by this last sentence is explained by the feeling that beliefs which one ought not to hold ought not count as knowledge. But this is to ignore the lesson of the envatted brain. Knowledge is not a simple matter of epistemic responsibility—if it were, the envatted brain would know lots of things about the external world. Knowledge is much more a matter of how one's beliefs are hooked up to the world. This fact lies at the root of the envatted brain's lack of knowledge of matters worldly. Once it is accepted, space is opened up for the conceptual possibility of the presence of the right sort of belief-world hook-up in the absence of proper epistemic responsibility. Once this conceptual possibility is admitted, the jarring note which accompanies the claim that Maud knows, even though she ought not believe, should disappear. In brief, accepting that knowledge requires the right kind of hook-up between world and belief removes the apparent implausibility attendant on the claim that Maud knows, though she ought not believe.

Bonjour presents a further objection to reliabilism, granting, for the sake of argument, in the Norman case that Norman does know that the President is in New York. He then presents what he takes to be an embarrassing consideration for the reliabilist. Norman comes to believe, on the basis of normal evidence, that the Attorney General is in Chicago. The evidence, however, is not strong enough to meet the requirement for knowledge. However, if Norman were asked to bet money on one of his beliefs being true, it would be more reasonable to gamble on the Attorney General's whereabouts, as he has no evidential basis for his belief about the President's whereabouts. Since the reliabilist claims Norman's belief about the President is knowledge, and since his belief about the Attorney General is not, the reliabilist is forced to the counterintuitive position that it is more rational to act on a mere reasonable belief than it is to act on a belief which is an instance of knowledge.

Two responses are open to the reliabilist. Firstly, if we examine the details of the case, it does not appear possible to both bet against the knowledge claim and persist in believing it. The reason for doubting the possibility of this case is more clear cut than the general reason given above. It has to do with the impossibility of holding contradictory occurrent beliefs. If Norman were asked to bet on either the whereabouts of the Attorney General or the President, and the case were to go as Bonjour describes it, presumably Norman would have to decide in which belief he was more confident. This, if he were being rational, could presumably only be done by examining his epistemic reasons for those beliefs. So, Norman would at some point arrive at a thought something like "Wait a minute, I can't bet on the President's whereabouts, I haven't got the faintest idea why I believe that in the first place." This thought, should, if he is being rational, lead him to reject (or at least suspend) his belief about the President's location.

Hence, if Norman were asked to bet on one belief or other, he would, if being rational, reject the belief about the President's whereabouts. But, if he did that, neither the reliabilist nor anyone else would be committed to saying that he knew this. The reliabilist can conclude that Norman does in fact know the whereabouts of the President, but if he were asked to bet on it, he would stop believing that the President is in New York , and hence, stop knowing it.

However, even if the above response is rejected, nothing about the Norman case need worry the reliabilist. For the case, so far at least, is underdescribed. The first question that the reliabilist ought to ask is whether Norman's belief about the Attorney General's location was formed by a reliable process. If it was, then it too, no less than his belief about the President's location, is knowledge. In this case, we do not have Norman haplessly betting on mere reasonable belief over knowledge; we have, instead, a case of Norman betting on one instance of knowledge over another.

On the other hand, if Norman's belief about the Attorney General's whereabouts was not formed by a reliable process the reliabilist must indeed swallow Bonjour's conclusion. In this case, we do indeed have Norman's betting on reasonable belief over knowledge—but this is not so counterintuitive as first appears. Recall that reliabilism is not supposed to be criterial for knowledge. Norman can know the President's whereabouts without knowing that he knows. He can, and in fact does, know without even believing that he knows. On the other hand, Norman believes that (it is probable that) he knows the Attorney General's whereabouts. It is more rational for him to bet on the whereabouts of the Attorney General, but this is not a problem for reliabilism. Reliabilism explicitly disavows the task of providing a way of telling, from the inside, when you know. It attempts instead to give an account of what conditions must be satisfied for one to be in the knower relation, regardless of whether one can tell, from the inside, whether these conditions are satisfied. It therefore leaves room for cases like Norman's. If the KK thesis were true, there would be something seemingly absurd about choosing reasonable belief over knowledge. But once that thesis is rejected, we do not have a case where Norman chooses reasonable belief over knowledge-under-that-description. As far as Norman is concerned, he chooses one belief over another, and in that, there is nothing problematic.

A related point here, which helps take the sting out of the conclusion the reliabilist must swallow, is that Bonjour's attack relies on the idea that the purpose of knowledge is to control and regulate behavior. Surely, we are encouraged to believe, if one has two different propositions in one's belief-box, such that one of them is known, whereas the other is merely believed, then one ought, rationally, be more inclined to act on the proposition which is known. But it is not the job of knowledge thus to regulate behavior. It is <u>belief</u> which has this function, not knowledge. One cannot, from the inside, tell the difference between a belief which is actually knowledge, and a belief of which one is merely convinced, and which may be erroneous.[28] One cannot, therefore, sort into a group all the things one knows, and use these to guide behavior. One can

[28]A caveat is here in order. This point does not hold of (at least a large portion of) a priori knowledge. Logical and mathematical knowledge for which one possesses sound proofs can, properly, be seen as knowledge, even from a first-person perspective.

only act according to what one believes. It is not at all obvious, therefore, that the conclusion that it is wiser to act, in certain cases, on mere reasonable belief over knowledge, is unwelcome.

To recap, it has been argued that cases like Maud's and Norman's need not trouble the reliabilist. Firstly, there are reasons for doubting whether the cases are genuinely possible. Even if we allow such worries to pass, we can still see the beliefs as being justified, even though the believers have no epistemic access to the justification. The reliabilist can allow that the agents ought not to believe, but in fact know the relevant propositions. Finally, the reliabilist can answer Bonjour's point about the irrationality of acting on mere belief over knowledge by remarking that there is simply no way to distinguish, from the inside, which beliefs we in fact know, from those of which we are merely confident.

### 1.3.2

### The Generality Problem

Richard Feldman,[29] and Earl Conee and Richard Feldman[30] have argued that the generality problem effectively refutes reliabilism. Feldman presents the generality problem for reliabilism as follows. According to reliabilism, a true belief counts as knowledge iff that belief was produced (or sustained) in a subject by a (relevant) reliable belief-forming process. Obviously an account of what is meant by 'relevant

[29] Richard Feldman, "Reliability and Justification", The Monist (1985): 159-173.

[30] Conee, Earl, and Feldman, Richard, "The Generality Problem for Reliabilism," in The Theory of Knowledge, ed. Louis Pojman (CA: Wadsworth, (1999), 343-357.

types of belief-forming processes' is required. Without such an account, as Feldman notes, we have no idea how to evaluate the theory, as we have no idea what its consequences might be.

However, Feldman argues, there seems no way to give an account of a process type without falling foul of one of two dangers, either of which is fatal to the theory. If the relevant types are characterized too narrowly, the relevant type for a given token may have only one instance. Should that token yield a true belief, the relevant type will be completely reliable, and the belief will automatically be justified. Every true belief will, trivially, become knowledge. Parallel results follow for false beliefs. This is the "Single-Case Problem."

On the other hand, specifying types too broadly leads to the "No-Distinction Problem." If the relevant process-type leading to perceptual beliefs were simply the process of using the visual system, for example, then the reliabilist would be committed to saying that a belief formed about a medium-sized physical object, in good light, from four feet away was exactly as justified as a belief formed in poor light, about a small object, from thirty yards. Clearly this is an unacceptable consequence.

The generality problem for reliabilism, then, is to specify what is to count as a process so as to avoid both of these problems. In this section, I shall first discuss the worries raised in the Feldman paper, using this discussion to develop a reliabilist solution to the generality problem. Following this, I will argue that this solution also avoids the problems raised by Conee and Feldman for reliabilist solutions to the generality problem. For the most part, the discussion will focus on visual beliefs. However, it is expected that the account yielded will generalize to cover other belief-independent processes. Where this seems dubious, the route to such a generalization will be sketched.

Reliable belief-independent processes can be thought of as functions which take as inputs information about distal states of affairs, and yield as outputs beliefs which are (usually) true.[31] Standard examples would be vision and (to a greater or lesser degree) the other sensory modalities. However, if we take vision as a single process, the no-distinction problem arises. One obvious suggestion is to slice belief-independent processes more narrowly than (e.g.) the visual system. Seeing from afar, then, will count as a different process than seeing from close up. In general, the specification of process types will include some reference to observation conditions. The account of a belief-independent process (BIP) will then look something like the following.

BIP: (1) A BIP is a function which takes as inputs information about distal states of affairs, and produces beliefs as output.

(2) Because information about distal affairs is relative to observation conditions, specification of a BIP must make reference to the observation conditions under which the belief was formed.

To this kind of move, Feldman (1985, 164) objects that, no matter how observation conditions are specified, numerous visual beliefs can be specified in the same observation conditions, and some of these will be more justified than others. As an

[31] See Alvin Goldman, "What is Justified Belief?" in Justification and Knowledge, ed. G. Pappas (Dordrecht: Reidel, 1979): 1-23.

example, Feldman notes the case where, seeing a distant mountain-goat, I form two beliefs—one being that there is a goat over there, and the other being that there is an animal over there. Intuitively, we would say that the latter belief is more justified than the former. The above definition of BIP, however, implies that the two beliefs are equally justified.

Feldman (1985, p.164) notes that we can avoid such problems by specifying in the account of the visual process, not only observation conditions, but also types of beliefs. Individuating process types along these lines, then, we add (3) to BIP.

> (3) Because the type of information which is relevant to the production of a belief is relative to the type of belief produced, specification of a BIP must make reference to the type of belief produced.

Feldman rejects such a proposal for two reasons. One is that he thinks that specifying what counts as the same kind of belief presents a problem no less challenging than the generality problem itself; the other is an objection based on the existence of borderline cases.

To motivate a method of individuating process types which escapes both of Feldman's criticisms, consider the following case. Billy is a reliable spotter of animals. Pretty much every time Billy says that an object on a distant mountain-top is an animal, it turns out to be just that. The same is true of situations where Billy gets a close-up look at the object. Billy, on the other hand is a lousy spotter of goats. In fact, whether near to or far from the stimulus, about half the time Billy says that the animal he can see is a goat, it turns out to be a sheep, and vice-versa. Imagine that you have been given a large sum of money to gamble on Billy's predictions. If you are rational, you will bet only on Billy's animal-predictions, keeping your money safely in your pocket when Billy makes a sheep-prediction, or a goat-prediction. Why is Billy to be trusted with one claim, but not the other? Clearly the answer lies in his respective reliability in each case. Billy is a more reliable animal spotter than goat spotter. One plausible explanation for this is that a different process underwrites his different kinds of beliefs. This lends credence to the suggestion of relativizing BIP to kinds of beliefs. The question is, can this relativization be spelt out in a principled way?

As Billy's visual beliefs are reliable or not depending on what property he is attributing to what object, we might at first attempt to specify the process in such a way as to make reference to the external object. Thus the process by which Billy comes to believe that there is an animal over there is the process of actually seeing an animal over there. This won't do. Firstly, in the case under consideration, it would guarantee the truth of Billy's belief as a matter of logic. If Billy comes to believe that that's an animal over there by seeing an animal over there, there is no possibility of error. But this is not because of any particularly reliable process. Rather it is because, unless there is an animal over there, we don't count his belief as formed by the same process, and so the reliability of the process is trivially guaranteed. Secondly, if the belief-forming process were specified in this way, Billy would count, at least some of the time, as a highly reliable spotter of goats. When is he a highly reliable spotter of goats? Why, when he forms his belief that a distant animal is a goat via the process of seeing a goat over there. But of course, Billy is not a reliable spotter of goats.

What is wrong with the present proposal is that it specifies a process by building into it the truth condition for the belief it produces. There is, however, something right about this proposal. It recognizes the fact that part of what makes somebody's visual belief reliable is something about the nature of the object the belief is about. If we could specify processes in such a way as to retain this advantage, while avoiding the triviality trap, we would be well on our way to the principled relativization of processes to beliefs that we desire.

Fortunately, there is a way to do just this. We need only borrow Goldman's (1999) account of "perceptual equivalents." Goldman's account is complex, and we need not go into all the details here. The basic idea is that two situations are perceptual equivalents (PE's) for one another if they would produce perceptual experiences which the perceiving subject would be unable to discriminate. Qualitative identity of percepts is not required for similarity here. Instead, two similar situations are PE's if the percept produced by one situation would not differ from the percept produced by the other, in any respect that would be causally relevant to the belief (Goldman, 1999, 175). Goldman (1999, 175) illustrates what is meant here by means of the following case. Trudy and Judy are twins. Sam's Trudy-produced percept differs from his Judy-produced one only in that the shape of Trudy's eyebrows differs from that of Judy's. Sam, however, does not (consciously) register this difference. The two percepts are thus PE's for one another, since they do not differ in any respect which is causally relevant to the belief.

Clearly, then, situations are PE's for one another, not necessarily in virtue of properties they actually exemplify, but in virtue of properties they are taken to exemplify (Goldman, 1999, 175). Equally clearly, putting things this way thereby relativizes our account of process types to persons.

We now add (4) to BIP:

(4) Two beliefs are of the same type iff, with respect to the property ascribed by the belief, the (similar) situations they are about are perceptual equivalents of one another.

The answer to Feldman's question of when two processes are of the same type, then, for perceptual beliefs, is that two beliefs are formed by the same process iff they are formed under the same observation conditions, are produced by similar perceptual experiences, and the property ascribed in the beliefs is (for the perceiver) similar.

It might be thought that relativizing to PE's would prevent us from generalizing the account of process types to non-visual belief-independent processes. In fact, all that is required for such a generalization is that there be situations which are (sometimes) indistinguishable from other situations for each sensory modality. Such situations would count as equivalents in the given modality. An example would be that of being unable to distinguish the voice of a friend from that of her sister on the telephone.

Relativizing to PE's enables us to explain how BIP avoids the Single Case Problem. What makes two situations PE's for one another is not properties the situations actually exemplify, but properties believed to be exemplified. And since, ex hypothesi, Billy regularly believes of situations in which the property of being a sheep is exemplified that they are situations in which the property of being a goat is exemplified (and vice versa), these situations are, for Billy, PE's for one another. Our account thus has it that it is the same process which produces both Billy's goat beliefs and his sheep beliefs. Suppose, however, that some unique scene has no actual PE's. Won't the foregoing account yield the consequence that a process which takes such a scene as input, and yields a true belief, will be completely reliable? If so, isn't that a bad result, of the type Feldman was worried about? The key to answering this point lies in the applicability of the predicate 'reliable'. This is a dispositional predicate, and it thereby carries implications about what would happen if the process were to be repeated. This suggests considering close counterfactuals to evaluate the reliability of processes which have unique scenes for inputs. A scene with no actual PE's will still have counterfactual PE's. If, when we go counterfactual, the process yields a high ratio of truths over falsities, it is reliable. If not, not.

BIP also avoids the No Distinction Problem. It delivers the correct verdict that Billy does not know that the animal in front of him is a goat (or a sheep), though about half of the time he guesses right. He does not know the animal is a goat (sheep) because he would still believe it to be a goat (sheep) if it were a non-goat (non-sheep) PE in front of him. By way of contrast, BIP allows, again correctly, that he does know that the object is an animal, because there are, for Billy, no non-animal PE's of animals. He would not, therefore, believe that there was an animal in front of him unless there were actually an animal in front of him. The process which produces Billy's animal beliefs is extremely reliable.[32]

So, contra Feldman, we now have an account of process types relativized to kinds of belief, which avoids the No Distinction Problem, and, by virtue of including the notion of perceptual equivalents, also avoids the Single Case Problem. Feldman (1985, 165) objects that any account of process types, relativized to belief types, will fail because of borderline cases:

> A person might view a succession of objects under exactly the same observation conditions and form beliefs about whether the objects have or lack a certain property.... Some of the objects the person sees may clearly have the property in question while others do not. As a result, some beliefs to the effect that the object has the property may be better justified than others.

They ought not to be, however, because they are all outcomes of the same process in the same observation conditions.

This objection has no force against the present proposal, however. To see this, consider whether seeing a clear example of an F which is G ought to count as a token of the same process type as seeing an example of an F which may or may not be G. Imagine you are an umpire in a tennis match, and you see a ball fall five feet outside of the court and thereby form the belief that the ball was out. Is the process which led to

[32] Of course, if there were a lot of cunningly crafted animal-looking robots in the vicinity, perhaps animals would have PE's for Billy. This does not constitute a problem for the present proposal, however. For in such a case, Billy would not be a reliable spotter of animals. Indeed, if the robots were manufactured well enough, it might never be discovered that Billy was not a reliable spotter of animals. This, of course, would be relevant only for whether Billy (or someone else) would know that he was a reliable spotter of animals. For a theory such as reliabilism, which drives a wedge between knowing and knowing that one knows, what is relevant is whether, in fact, Billy is reliable. Bearing this in mind, it does not seem untoward to say that whether or not Billy is reliable depends, among other things, on whether or not there are cunningly crafted animal-looking robots in the vicinity.

that belief the same process as that that occurs when you see a ball land right by the line (so that it is unclear whether the ball is in or out)? If we specify processes in such a way as to make reference to PE's, the answer is clearly no. The two cases are PE's for practically no-one. Therefore, the beliefs in each case are formed by different processes.

Borderline cases are, of course, PE's for each other (in some cases at least). This, far from being an objection to the present theory, is a welcome consequence. For, since borderline cases are sometimes PE's for each other, there are many cases where our belief that an F was G would have been caused by a borderline F which was not G, although it appeared to be. Which is to say that our beliefs about borderline cases are not particularly reliable. Which is in turn to say, properly, that such beliefs ought not count as knowledge.

As BIP avoids both the Single Case Problem and the No Distinction Problem, specifies how to individuate kinds of beliefs, and accounts for borderline cases, I take Feldman's original criticisms to have been met. I turn now to the Conee and Feldman paper. The authors of that paper classify reliabilist solutions to the generality problem into six types, and argue that each solution-type faces insurmountable difficulties. Of these types, BIP is closest to the "Maximum specificity and Narrow causal types" category. The terminology originates from Alston,[33] who, like me, explicates processes in terms of functions. For Alston, a function is maximally specific when "any difference in input that is registered by the function indicates a different function" (Alston, 1995, 26). This idea fits well with my use of PE's. When two scenes are PE's for each other,

---

[33] William Alston, "How to Think About Reliability," Philosophical Topics (spring 1995): 1-29.

no difference in input is registered by the function, and so, the process which takes either scene as input, and forms a belief of a given type, will be the same process in each case. There is, however, one difference between my account of process types and the Maximum Specificity proposal. To see this, consider Conee and Feldman's (1999, 350) gloss on Alston's suggestion:

> The maximum specificity proposal is the idea that the relevant type includes all and only process tokens with the same causal features: they all begin with experiences with the same causally active features, are followed by subsequent events with the same causal features, and have the same belief as output.

This, they note, leads to the unacceptable conclusion that no process can produce more than one belief content. Call this the "Single Content Problem." BIP avoids this problem. It counts different beliefs as being produced by the same process type if the scenes which lead to those beliefs are, relative to an observer, PE's for each other, and if the properties ascribed in those beliefs are similar. In our example, the process which led to Billy's belief 'There is a sheep' is the very same process which led to his belief 'There is a goat'. Since sheep-scenes and goat-scenes are PE's for Billy, no difference in input will be registered by the function, and so the process which takes either scene as input and forms either belief, will be the same regardless of which scene is inputted, and regardless of which of the two beliefs is outputted.

In spite of this difference between Alston's proposal and my own, BIP is closer in spirit to the Maximum Specificity proposal than to any other proposal considered by Conee and Feldman. If, then, BIP avoids the other problems they raise for that solution, I take it to provided an acceptable reliabilist response to the generality problem.

Conee and Feldman raise a follow-up objection to the Single Content Problem which might be thought to apply to BIP, though it in fact does not. The objection targets versions of reliabilism according to which processes can output only one belief content. If such a process led to a belief in a necessary truth, then the reliability of that process is guaranteed by the mere content of that belief, since the belief produced could never be false. Thus, the nature of the process itself is immaterial. But the appeal of reliabilism stems from the conviction that what makes some beliefs knowledge has something to do with the nature of the process which produced it. Thus this consequence is at odds with the spirit of the theory.

Since BIP does not slice processes so thin that they can have only one output content, the objection does not immediately apply. But let us attempt to alter the case so that it does. Suppose some process takes as input some scene, or any of its PE's, and outputs a belief in one of a range of necessary truths. [34] In altering the case so that it can apply to BIP, however, we have removed the very feature which made it problematic to some versions of reliabilism. For, if the same process can produce belief in any one of a range of necessary truths, what is left of the claim that the nature of the process is rendered irrelevant by the mere content of the belief? Surely if the process can produce

[34] Even as modified, this is not yet a counterexample to BIP. BIP types its tokens in part on the basis of similarity between beliefs, and Clause (4) has it that in order for two perceptual beliefs to be relevantly similar, those beliefs must be about perceptual scenes. In order for the case to be a counterexample to BIP, the belief in necessary truth, p, would have to be (i) a necessary truth about the given perceptual scene, which is (ii) formed directly from the perceptual scene, without any mediating inference. Even granting that there are a posteriori truths which are necessary (and thus, that (i) can be satisfied) it hardly seems likely that such truths as 'Water is necessarily H2O' can be known without inference. Let us waive these points however, and suppose that some process takes a perceptual scene as input and yields a belief in one of a range of necessary truths (a posteriori or a priori). Perhaps this process works by magic, or perhaps it is akin to clairvoyance.

belief in a variety of truths, the explanation for this reliability over a range of contents would have something to do with the nature of the process itself. [35]

Conee and Feldman's next objection is that maximally specific solutions cannot distinguish the epistemic status of lucky guesses from that of expert judgements. It may happen that I always correctly associate some particular percept with gingko trees, not because of any aptitude in recognizing ginkgo trees, but in virtue of a lucky guess, or for some similarly poor reason. Since this would be a case of a process leading from a distal scene to a true belief, BIP would count it as justified, and thus confer the same status on lucky guesses as it would on expert judgement.

In response to this, the reliabilist should ask whether the process in question counterfactually yields a high ratio of truths over falsities. If it does not, then BIP will not count it reliable, and thus not deliver the verdict that the ginkgo belief is justified. Since lucky guesses are unlikely to be counterfactually reliable, it is correspondingly unlikely that BIP will place them on the same epistemic footing as expert judgement. If, however, the process is counterfactually reliable, then BIP will hold that the process yields justified beliefs. This, in essence, is Conee and Feldman's final objection. It may be that my ginkgo beliefs are somehow counterfactually reliable, even though I am unable to offer any justificatory reason for why I have the beliefs I do. The reliabilist is forced to hold that my gingko beliefs are knowledge. But if the process which produces my ginkgo belief is that of "applying some ridiculous and unjustified sort of

[35] One might here object that in order for a process to yield a belief in a necessary truth, having taken only a perceptual scene as input, it would have to be a bizarre process like clairvoyance or magic. Thus, it could not yield knowledge, and so BIP gives the wrong verdict here. I will give my answer to this kind of objection in my discussion of Conee and Feldman's use of Bonjour-style counterexamples.

numerology to the topic"(Conee & Feldman, 350), then it seems this is the wrong verdict.

Strictly speaking, this is not an objection to BIP. BIP is an account of non-inferential belief-forming processes, and the current case, as described, involves an inference between percept and belief. Let us grant, however, that the case can be altered so as to deal with this objection. The first thing to note is that this case is structurally parallel to other cases in the literature, most notably Bonjour's clairvoyance cases.[36] As we have seen, Bonjour's clairvoyant, inexplicably, but with nomic reliability, forms beliefs about the President's whereabouts. She has no reason for believing what she believes, though her genuine clairvoyance subconsciously delivers reliably true beliefs. Similarly here. I, ex hypothesi, am a nomically reliable ginkgo tree-spotter. I can offer no reason for why it is that I believe the tree in front of me to be a ginkgo, but in all close possible worlds, when I say it's a ginkgo, a ginkgo it is. Some actually operative but subconscious process delivers the correct verdict every time, in spite of my inability to verbalize it. (I note in passing that neither Bonjour's nor Conee and Feldman's cases specify the case in this way. But if the process yielding the belief is not psychologically or neurologically instantiated, then the cases involve some very dubious metaphysics, and their intuitive appeal is very greatly reduced. It is hard to see how a process which is not only not consciously accessible, but also neither neurologically nor psychologically instantiated, could lead to any beliefs. At the very least, the reliabilist can claim that what has been offered as an objection to reliabilism is, at best, a sketch of

[36] My source for this point is Michael Levin and Jonathan Adler "Is the Generality Problem Too General?" Philosophy and Phenomenological Research, LXV (2002): 87-97.

an objection, one that requires no answer until that sketch is completed).

In any event, the reliabilist response to Bonjour, and to Conee and Feldman, is that such cases derive their plausibility (as counterexamples to reliabilism) by illicitly importing intuitions about the real world into the worlds described by the cases. Our intuition that Bonjour's clairvoyant does not know is shaped by our beliefs about the impossibility of clairvoyance in our world. But if Maude is clairvoyant, and always reliably and correctly reports the President's whereabouts, why not say she knows? What reason, other than our beliefs about the impossibility of clairvoyance—which are ruled out by the description of the case—is there to deny that Maud knows? Similarly, if I am a nomically reliable spotter of ginkgoes, what reason—other than our real-world intuition that I could not have such an ability—is there to say that I don't know that there is a ginkgo in front of me? Certainly, if you knew of my ability, you would take yourself to know that something was a ginkgo when I said it was. You would bet with me rather than against me on the issue. If my ability allows you to know, why doesn't it allow me to know? If the response is that you know the object is a ginkgo by induction on my past successes, why not say, with the reliabilist, that until I perform a similar induction, I know the object is a ginkgo, but don't know that I know?

I take it that my responses to the Single Content Problem and the problem raised by lucky guesses are unobjectionable. I have further argued that, as BIP avoids the Single Content Problem, it does not succumb to the objection from necessary truth. The response to the Bonjour-type case, will, to some, sound implausible. I have attempted to soften the blow, suggesting that such cases are question-begging unless they specify what is wrong with the reliabilist intuition that true beliefs formed via a process which is nomically reliable are knowledge—even if that process is inaccessible to the believer's consciousness. Without such a specification, Bonjour-type cases are harmless to reliabilism. And, if the threat posed by the generality problem ultimately rests on such cases, it is no more life-threatening to reliabilism than are they.

We are left with the question of how belief-dependent processes fare in the face of the generality problem. Fortunately for the reliabilist, the generality issue here is much more easily dealt with. As Levin[37] remarks, reliabilism can take different modes of inference as different processes, with differing degrees of reliability. Reasoning via modus ponens, since it is infallible, is a highly reliable process. The same goes for all valid logical rules of inference, and for mathematical rules.[38] Concluding p from the mere fact that the pope has asserted p is (on the assumption that the pope is, after all, fallible), less reliable. The reliabilist, then, can take advantage of these ready-made distinctions between reasoning processes and use them to defuse the generality problem for belief-dependent processes. I take it that the work of cognitive psychologists in enumerating and evaluating human problem-solving methods will be relevant to a full account of which processes are reliable. To say this, however, is not to say that the generality problem has not been answered. It is, rather, just to admit what there is no reason to deny, that empirical work in psychology will have something to contribute to a fuller understanding of human cognition.

---

[37] Michael Levin, "Reliabilism and Induction," Synthese 97 (1993): 300.

[38] The issue of the reliability of our non-inferential mathematical beliefs will be taken up in Chapter Five. Questions examined there include whether intuition is the process by which we come to have these beliefs, and, if so, whether this is problematic.

### 1.3.3

### The Reliabilist Account of Warrant: Redundant or Just Plain Wrong?

Evidentialism is the thesis that we should believe all and only those propositions for which we have sufficient evidence. Proponents of this thesis are fond of claiming that, since theirs is obviously the correct notion of warrant, reliabilism is either redundant (if it yields the same account of warrant) or mistaken (if it yields a different account).

Feldman and Conee[39] make use of Bonjour cases to argue as follows:

1. The reliabilist can respond to Bonjour's cases in one of only two ways. They can claim either (a) Norman's beliefs are well-founded,[40] or (b) although Norman's beliefs are not well-founded, reliabilism does not yield the consequence that they are.
2. Response (a) is implausible (Feldman & Conee, 1985, 28).
3. The only way the reliabilist can plausibly opt for response (b) is "to specify the relevant types of belief-forming processes in evidentialist terms" (Feldman & Conee, 1985, 30).
4. This leads to a "roundabout approximation of the straightforward evidentialist thesis" (Feldman & Conee, 1985, 30).
5. Reliabilism, then, is either false or it is evidentialism in reliabilist clothing.

---

[39] Richard Feldman and Earl Conee, "Evidentialism," Philosophical Studies, 48 (1985): 15-34.

[40] Well-foundedness is an evidentialist notion, to the effect that S's belief that P is well-founded iff S has evidence for P, and S has no more inclusive body of evidence which implies not-P.

My response to the Bonjour cases discussed above should indicate that I do not think this is the correct analysis of the situation. I reject step (1) of the evidentialist argument. From a reliabilist perspective, making well-foundedness the crucial epistemic concept in evaluating the Bonjour cases is illegitimate. Well-foundedness is explained in terms of evidence in favor of a belief, which evidence is possessed by the subject. Clearly, the subjects in the Bonjour cases lack evidence for their beliefs. But to use this point against the reliabilist is simply to beg the question. The question is whether, given that the beliefs are held in the absence of evidence, but are yet reliably produced, the subjects can properly be said to know. To decide this question by invoking a principle which refuses to countenance belief produced otherwise than via evidence is clearly illegitimate. Such a move cannot show that reliabilism is either false or redundant.

A further consideration which tends to undermine Feldman and Conee's conclusion is that the reliabilist can offer an account of why taking consideration of evidence is desirable. That account is simply that considering relevant evidence is a reliable method for generating true beliefs. This is why we care about evidence at all. If reasoning according to the evidence were not truth-indicative in general, there would be no rationale for evidentialism at all. I submit that the evidentialist can offer no compelling reason for belief in his thesis that does not ultimately fall back on reliabilist terms—viz. the reliability of evidence. It is not that what is correct about reliabilism is already captured by evidentialism, it is rather the other way around.

Haack offers a different argument for the conclusion that reliabilism must ultimately fall back on evidentialism.[41] She (1993, 143, 150, 210) notes that, in deciding which processes we take to be reliable, we use our evidence as to their reliability as our criterion. Because the reliabilist is unable to step outside his beliefs about reliability and grasp the reliable-in-itself, it follows that all he can do is use what evidence he has at his disposal in order to form beliefs about which processes are reliable. In falling back on evidence to make decisions about reliability, Haack claims, the reliabilist becomes a default evidentialist.

Michael Levin[42] offers the following, to my mind convincing, response:

> All we can ever do is ascribe a predicate to what we think it applies to. However, this does not impart an evidentialist spin to all predicates. Take "binary," which astronomers use of stars with companions. Naturally, astronomers will call a star "binary" only when they believe it has a companion, but they are not thereby committed to a definition of "binary star" as one believed to have a companion. ... Likewise, the fact that a reliabilist will rate a claim justified only when he believes its generating process is reliable does not commit him to an evidentialist definition of justification that refers to his or anyone else's beliefs.

In brief, although we will have various beliefs about which processes are reliable, only those true beliefs that are in fact reliably formed are genuine instances of knowledge.

Another attempt to argue for the conclusion that reliabilism, if true, is redundant goes as follows. The reliabilist acknowledges that our best policy is to base our beliefs on evidence. He thereby acknowledges that we ought to attempt to have our beliefs

---

[41] Susan Haack, Evidence and Inquiry: Towards Reconstruction in Epistemology (Cambridge, Massachusetts: Blackwell, 1993). See also John Pollock, "Epistemic Norms," Synthese 71 (1987): 61-95.

satisfy evidentialist criteria. But beliefs which satisfy evidentialist criteria are thereby warranted. Since such beliefs are warranted independently of reliability considerations, reliabilism is therefore redundant.[43]

The plausibility of this argument depends on a tendentious understanding of what is meant by 'warranted'. If 'warranted' means something like 'assertible, given the available evidence', then for sure the reliable connection is redundant. It hardly needs adding that such a reading of 'warrant' is, in the current context, completely illegitimate. On the other hand, if warrant is understood instead as 'whatever must be added to true belief in order to have an instance of knowledge',[44] the objection has no force. For we can have many evidentially warranted beliefs that fail to be knowledge, as is made clear by Gettier cases. And this fact shows that the reliable connection is not made redundant by a belief satisfying evidentialist criteria. Once more reliabilism is shown not to be a roundabout version of evidentialism.

So far we have accepted the evidentialist line that we should believe all and only those beliefs for which we have sufficient evidence. Instead of disputing this thesis, we have shown that its appeal is explained via the reliable connection between evidence and truth. We have, that is, argued that what is correct about evidentialism must ultimately be explained in reliabilist terms, so that it is evidentialism, and not reliabilism which is redundant as an account of warrant. But things are worse than this

---

[42]Michael Levin, "You Can Always Count On Reliabilism," Philosophy and Phenomenological Research, Vol. LVII, No. 3 (1997): 3.

[43] Jonathan Adler raised this objection in connection with an unpublished paper of mine.

[44] This is how Alvin Plantinga suggests we understand it, so as to avoid prejudging issues such as the one under discussion in the text. See his Warrant: The Current Debate (New York: Oxford University Press, 1993a): 3.

for the evidentialist. For, even if we leave to one side the point that evidence is desirable because of its reliability, the evidentialist thesis cannot be the whole story about warrant. Considerations arising from mathematical epistemology serve to show this.

Evidence, Katz observes, has no place in mathematics.[45] The warranting role which evidence plays in empirical belief formation is played, in mathematics, by proof. If evidence did play a warranting role in mathematics, then the accumulation of instances which conform to Goldbach's conjecture, for instance, ought to have provided enough evidence for us to say that we know this conjecture to be true. There are confirming instances aplenty, but we cannot say that we know that Goldbach's conjecture is true. Why is this? Because proof, and not evidence, is what provides warrant for mathematical beliefs.[46]

Not only does evidence not play a role in mathematics, it cannot. On the traditional picture of mathematical truths, such truths are necessary. Evidence, however,

---

[45]In conversation. But, for similar remarks, see Katz (1998): 36-41.

[46] I am unsure whether Katz meant these points to apply only to the kind of warrant required to establish mathematical truths, or whether he meant them to rule out even the possibility of mathematical truths being known through induction by the mathematically unsophisticated. Katz's point seems undeniably true for Goldbach's conjecture, but seems troublesome if pressed into service to deny that the shopkeeper, ignorant of mathematical proof, is not warranted in his belief that fifty pence is the correct change to give when one pound is offered in payment for a fifty pence purchase. Such warrant would, of course, count for nought in a mathematics department, but the shopkeeper's belief seems, in some subordinate sense, to be warranted.

Another example of empirical warrant in mathematical knowledge is my coming to know the truth of a given theorem by being told by an honest mathematician that the theorem has been proven. My knowledge here would be a posteriori. However, this kind of empirical warrant also seems to be subordinate to a priori warrant. I could only have such knowledge provided that an a priori proof of the theorem were available. Thus this kind of a posteriori knowledge of mathematical truth requires the prior establishment of the result by a priori methods. So, perhaps it is best to say that, in the context of discovering or establishing mathematical truths, evidence plays no role in mathematics. Once a truth has been established, evidence can supply warrant for mathematical beliefs in the subordinate kinds of ways listed here. This more restricted point still suffices to make trouble for the evidentialist. I thank Michael Devitt for pushing me to be more explicit on this point.

is an empirical notion which is used to decide questions of how things actually are, based on what else obtains in the actual world. Evidence, therefore, cannot tell us which of the things which happen to be true are necessarily so.[47] But necessary truths are the subject matter of mathematics, so evidence is not an appropriate tool for mathematics.[48]

Now it may be that the evidentialist can make a terminological change in her theory in order to allow for the special nature of knowledge-gathering in mathematics. She may adopt some broader notion, such as "having appropriate reason," which encompasses both evidence and mathematical proof. The trick, however, is for her to describe such a broader notion of warrant in such a way which does not simply include mathematical proof because of its reliability. This would, of course, play right into the reliabilist's hands, allowing him once more to return the accusation of redundancy to the evidentialist. Considering, however, that the desirability of valid proof forms resides in the fact that the truth of the premisses gives us a completely reliable guarantee of the truth of the conclusion, it is rather difficult to see how to avoid appealing to reliability in an explanation of the warrant provided by such proofs.

I conclude that the reliabilist can not only avoid the evidentialist objection, but can also turn that objection back against the evidentialist. The correct answer to the

---

The relationship between proof and the a posteriori will be discussed in great detail in Chapter Three. For now, it is enough to note that Katz's point can be made by explicitly restricting the domain to the type of warrant required to satisfy mathematicians.

[47] Michael Devitt objected to an earlier draft of this section, on the grounds that we can have empirical knowledge of necessity. We know that water is necessarily H2O, and, he holds, we know this through empirical evidence. On this, I think he is mistaken. We surely know empirically that water is H2O, but we know that this identity is a necessary one not through empirical science, but through a priori philosophy. Granted this, my point in the text stands: evidence cannot tell us which of the things that happen to be true are true necessarily. Evidence thus cannot play a role in establishing mathematical truth.

question which heads this section is a resounding "Neither." And this answer completes our defense of reliabilism. The next step, to be undertaken in Chapter Two, is to defend the notion of a priori knowledge from Quine's attacks.

[48] Of course, for the evidentialist who does not accept that mathematical truths are necessary, this second objection does not arise.

# Chapter 2

## Quine's Attacks on the A Priori

### Introduction

The purpose of the remaining chapters is to defend the notion of mathematical knowledge as knowledge a priori. To that effect, this chapter will examine Quine's arguments against there being any a priori knowledge at all. The following chapter will then defend the conception of mathematical proof as being a priori from attacks due to Tymoczko and Kitcher. The final two chapters will present and defend a reliabilist account of a priori mathematical intuition.

Discussions of Quine's objections to a priori knowledge typically cite "Two Dogmas of Empiricism" as the work in which those objections are presented. This work contains two of Quine's arguments against a priori knowledge: his argument that there is no intelligible notion of analyticity, and his argument from holism. However, commentators as diverse as Gibson[49] and Bonjour[50] see at least one further argument elsewhere in Quine's corpus. Both Gibson and Bonjour see an argument against the a priori stemming from Quine's naturalized epistemology.

---

[49] Roger Gibson, Enlightened Empiricism (Tampa: University of South Florida Press, 1988).

[50] Lawrence Bonjour, In Defense of Pure Reason (Cambridge: Cambridge University Press (1998).

There are, then, three arguments of Quine's which will receive attention in this chapter. They are

1. The argument from the unintelligibility of analyticity.
2. The argument from holism.
3. The argument from naturalized epistemology.

Before examining these arguments, however, I wish to consider a point due to Laurence Bonjour (1998, 63 and 81), which, if good, would render superfluous any discussion of Quine's attacks on the a priori. Bonjour's point takes the form of a dilemma for Quine. Quine claims that there are no a priori truths. But what is the status of this claim? It must, Bonjour tells us, itself be either a priori or a posteriori. If the former, then Quine's position is self-defeating. But it cannot plausibly be construed as a posteriori. This is for two reasons. Firstly, appeals to experience have no bearing on whether or not a priori justification is possible. And secondly, it does not seem possible that the claim that there are no statements which are justified a priori could be shown to be true via experience. Bonjour is not explicit on why this is so, but presumably he has in mind something to do with the impossibility of observing all statements.

Quine himself provides the answer to Bonjour's point. He has written again and again[51] that all sentences are a posteriori. If we take this seriously, then the claim that there are no a priori truths is, like all other claims, an a posteriori one, itself open to empirical refutation. What then of Bonjour's claim that so construing Quine's claim

[51] See, for example, Quine (1961, 42). See also his "Reply to Vuillemin," in Hahn, L., and Schilpp, P., (1986), The Philosophy of W.V. Quine, ed. L.Hahn and P. Schilpp. La Salle, Ill.: Open Court. See also his "Comment on Quinton," in Barrett and Gibson (1990).

makes it unjustifiable? Well, if Quine's empirical holism can itself be defended, this would justify the claim that there are no a priori truths by virtue of rendering such truths explanatorily redundant. If justification is plausibly construed as a matter of a statement's position in a holistic web of purely empirical statements, there simply is no place left for the a priori. [52] There seems little wrong with so arguing. We are, of course, left with the rather large question of whether empirical holism is a plausible account of justification. This question will be examined in detail in section 2.2. But for now, the point is that, pace Bonjour, there is no quick dismissal of Quine's position as self-defeating.

## 2.1

**The Argument From the Unintelligibility of Analyticity.**

Bonjour (1998, 64) remarks on the striking absence of any explicit mention of the a priori in the first half of Quine's "Two Dogmas." An extremely unsympathetic critic of Quine, he (1998, 64) goes on to claim that it is "easy to show that Quine's grasp of the main concepts and distinctions in the area ... is far from sure." Accusing Quine of conflating the analytic/synthetic distinction with the a priori/a posteriori distinction, Bonjour (1998, 65) further asserts that:

> Such carelessness would be objectionable enough in any case, but it is all the more disconcerting here because Quine's major objection, at least to the concept

[52] Oddly enough, though there is no discussion of this possibility in his (1998), in an earlier work (The Structure of Empirical Knowledge, Cambridge, Massachusetts: Harvard University Press. 1985, (195)), Bonjour notes the likelihood of such a Quinean move. In that work, Bonjour counters this move with the claim that the fact that a sentence occupies a certain position in a web of sentences gives us no reason to believe that sentence to be true. The question whether Quine offers satisfactory notions of truth and justification will be examined in sections 2.3.1.1, and 2.3.1.2.

> of analyticity, is that it is unintelligible, an objection that is hard to take seriously when even minimal efforts at clarification have not been made.

Bonjour (1998, 65) next makes what he considers a "speculative diagnosis," to the effect that Quine's focus on analyticity might be due to his endorsing the view that if there were any a priori truths, they would be analytic. But this, far from being a speculative diagnosis, seems to be the standard reading of the Quine of "Two Dogmas." Orenstein,[53] Haack (1993, 122), and Katz (1998, 184), inter alia, all interpret Quine as assuming with the Logical Empiricists that there are no synthetic a priori truths. For one sharing this assumption, an attack on analyticity simply is an attack on the a priori.

At any rate, I intend to read Quine's attack on analyticity as being an attack on the a priori. The argument, then, is familiar. No clear definition of analyticity has been given. The failure of all attempts to define the concept, without using similarly ill-defined notions (such as synonymy or necessity) does not offer much hope of future success. If, as seems likely, any account of analyticity will have to use similarly ill-defined intensional concepts in the account, analyticity admits of no satisfactory definition.

This argument has not, to put it mildly, been universally well received. The main goal of this section is to examine how it stands up to a sample of the criticisms directed at it. Putnam, for example, objects to Quine's line of argument, remarking that the only evidence given for the claim that the notion of synonymy was hopelessly vague was that

[53] Alex Orenstein, Willlard Van Orman Quine (Boston: Twayne, 1998), 85-6.

Quine himself could not clarify the concept. Putnam then, accuses Quine of committing an ignoratio elenchi.[54]

But this criticism misses an important point, one brought to light by Gibson (1988, 96). Quine, in Word and Object,[55] claims that he would be satisfied with a rough characterization of the notion of analyticity (or any notion from its circle of terms) in behavioral terms. This, in turn, raises the question as to why the clarification of the distinction must be couched in behavioral terms. The answer, as Gibson points out, is that given Quine's naturalism, and its attendant behaviorism, behavioral evidence is the only evidence there is in linguistics. It follows that if a notion is to be clarified in an acceptable manner, then for Quine, this must be done in behavioral terms. If, then, we take Quine's behaviorism seriously, we no longer have an ignoratio elenchi. Quine's argument is that the only evidence there is is incapable of supporting a hard and fast distinction between statements which are analytic and statements which are synthetic.

One might reasonably protest at the request for behavioral evidence. After all, there will, it seems, be no difference in the behavior of speakers of a language between cases of analyticity and cases of obviously true, but mundane, synthetic sentences such as 'There have appeared to be red objects'. Both sorts of sentence would elicit universal assent from competent speakers.[56] So, in restricting the available evidence, Quine

[54] Hilary Putnam, "'Two Dogmas' Revisited," in Realism and Reason, Philosophical Papers, Volume 3 (Cambridge: Cambridge University Press, 1983), 88. Putnam was not the first to make this point. See also H.P. Grice and P.F. Strawson, "In Defense of a Dogma," in Readings in the Philosophy of Language; ed. J. Rosenberg and C. Travis (Englewood Cliffs, N.J.: Prentice Hall, 1971), 87.

[55] Quine, Word and Object(Cambridge, Mass.: MIT Press, 1960), 207.

[56] Quine (1960, 55) stops just short of saying this when presenting his own notion of stimulus-analytic: "A sentence is stimulus analytic for a subject if he would assent to it, or nothing, after every stimulation (within the modulus)." Morton White explicitly makes this point in his "The Analytic and the Synthetic, an Untenable

renders impossible an acceptable characterization of the notion of analyticity. Certainly, both Bonjour (1998, 69) and Gibson[57] agree that Quine's behaviorism limits, ahead of time, what can count as an acceptable answer to this question. This behaviorism is a consequence of Quine's naturalism, which, as Bonjour remarks is roughly the claim that there is no genuine knowledge outside of empirical science. It appears, then, that Quine's argument against analyticity escapes the charge of being an ignoratio elenchi only at the cost of begging the question against the a priori.

Katz,[58] however, cautions against dismissing Quine's argument here on the grounds of its behaviorist assumptions. Katz tells us (1990, 179) that Quine's behaviorism is not "the fierce reductive doctrine of days gone by." It is instead "merely a way of putting the study of language on a par with other sciences by requiring the linguist's theoretical constructions to be justified on the basis of objective evidence." The claim that Quine's behaviorism is not of an objectionable sort receives support from a letter from Quine to Gibson, in which Quine appears to distance himself from behaviorism. Quine writes:

> When I have stressed that language is learned through observation of overt behavior without telepathic aids, I have encapsulated the point by saying that linguistics has to be behavioristic; but if the term does not fit my account, the term is what should be dropped (in Gibson 1988, 129).

---

Dualism," in John Dewey: Philosopher of Science and Freedom, edited by Sydney Hook (New York: Dial, 1950): 325.

[57] Roger Gibson, The Philosophy of W.V.O. Quine, An Expository Essay (Tampa: University of South Florida Press, 1982): xx, and 96-100.

[58] Katz, The Metaphysics of Meaning (Cambridge, Mass.: Bradford MIT, 1990).

Putting together what has been said so far, Quine's argument against analyticity is not an ignoratio elenchi if his restriction of evidence to behavioristic evidence is acceptable. If that behaviorism amounts to no more than requiring objective evidence for the linguist's theoretical constructions, it surely is acceptable. Given this, the demand for behavioristic evidence begs no questions.

At this point, however, we are in a position to show that Quine's argument against analyticity does not go through, regardless of what one thinks about the restriction of evidence to behavioral evidence. If Gibson and Bonjour are correct, and Quine's behaviorism is of a sort which, from the outset, proscribes a priori knowledge, then his argument begs the question. If, on the other hand, Katz is correct, and Quine's behaviorism is not of a sort which proscribes a priori knowledge, then no question is begged. However, in this case Quine's argument is not strong enough to support his conclusion. Quine's argument becomes hostage to the possibility of a satisfactory account of analyticity being presented. Indeed, Katz's own account seems to be just such an account. Katz writes:

> There is, then the option of explaining meaning, synonymy, and analyticity on the model of Chomsky's explanation of syntactic notions like 'well-formed'. We can construct an abstract system of semantic representations that formally describes the meaning of sentences, characterize semantic notions like meaningfulness, synonymy, and analyticity in terms of such formal representations, and then justify both the representational system and the definitions indirectly on the basis of how well they predict and explain judgements of fluent speakers about such semantic properties and relations of sentences.[59]

[59] Katz, "Common Sense in Semantics," Notre Dame Journal of Formal Logic (April 1982): 192.

As Chalmers Clark remarks,[60] this approach conforms to Quine's demand for an acceptable definition of a linguistic notion to be subject to behavioral evidence. The behavioral evidence is provided by the judgments of fluent speakers. The abstract representations posited by the theory earn their keep by providing successful predictions and explanations of the judgments of fluent speakers—just as, in Quine's conception of the physical sciences, theoretical posits are acceptable only if they facilitate successful predictions.

On neither view of Quine's behaviorism, then, is there a successful argument from the unintelligibility of analyticity to the non-existence of the a priori.

## 2.2

### The Argument From Holism

In section six of "Two Dogmas," Quine presents his now famous metaphor of the web of belief. Our beliefs form an interconnected whole. It is this whole, rather than any individual belief, which faces the "tribunal of experience." Those beliefs which are closer to the periphery are more likely to be revised in the face of recalcitrant experience than are those closer to the centre, but in principle, even the most central beliefs may be revised. Quine illustrates this point by noting the conceivability of a revision of the law of excluded middle in order to accommodate results from quantum physics.

---

[60] Chalmers Clark, (1998, 283), "The Art of Science: Quine and the Speculative Reach of Philosophy in Natural Science," in Dialectica, Vol. 2, No.4 (1998): 283.

Given an understanding of a belief which is justified a priori as being one which is immune to revision for empirical reasons, it follows that there are no beliefs which are justified a priori. The felt certainty of basic logical and mathematical beliefs is then explained on the basis of those beliefs being very close to the center of the web, and therefore being less likely to be revised in the face of recalcitrant experience. The effect of the argument from empirical holism, then, is to show that (1) it is a mistake to think that there are any a priori truths, and that (2) we can account for our confidence in the truth of basic mathematical and logical statements without assuming these to be a priori (or necessary), on the basis of their position in the web.

If Quine's argument is good, it seems that rationalists have had the rug pulled from under their feet. Answering Quine on this point is, then, a matter of some urgency for the defender of a priori knowledge.

Since we are addressing an argument which moves from the claim that holism is true, to the conclusion that there is no need for the notion of a priori knowledge, the best way to approach this argument is to ask how its main premiss is supported. How is holism argued for? The main consideration marshaled in favor of holism is the Duhem thesis (Quine 1961, 41; Orenstein 1977, 85-6).[61] Gibson (1988, 33) calls this the scientific practices argument. As a matter of empirical fact, a scientist who is testing some claim H must assume some auxiliary assumptions, A. If the experimental result threatens the truth of H, H can always be saved by making drastic enough changes in some or all of the claims in A.

---

[61] See also Quine's "Reply to Vuillemin," in Hahn and Schilpp, 1986, 619. The point is also made in the authors' introduction to Barrett and Gibson 1990, 10-11.

As a point about the rational revisability of empirical sentences, this seems undeniable. As both Katz (1998, 50) and Bonjour (1998, 76) have argued, however, no argument which depends on interpreting the Duhem thesis as a thesis about the rational revisability of all statements can possibly be an argument against a priori truth. At least not if, by an a priori truth, we mean one not vulnerable to empirical refutation. Without a question-begging assumption that a priori truths can rationally be given up in the face of recalcitrant experience, holism cannot be established via Duhem. A believer in a priori truth might, for instance, hold a logical law, like noncontradiction to be such a truth. If by this he means that the status of this truth does not depend in any way on empirical considerations, and that it is not vulnerable to being undermined by experience, then it is no argument against him to simply claim that the lesson to be drawn from the Duhem thesis is that all statements are open to possible revision in order to save some experiential H. This is simply to deny (without argument) that there are any a priori truths. The apriorist can accept the point that, in the face of recalcitrant experience, giving up some logical law, used in the derivation of H from (O+A) would save us from the undesirable conclusion not-H were it epistemically acceptable to give up an a priori principle in the face of experience. It is part of his position, however, that it is not epistemically acceptable to give up an a priori law for experiential reasons. To defeat the apriorist here, in addition to the Duhem thesis, one would need an argument to show the rational revisability of a priori propositions for empirical reasons.

One might claim that Quine has supplied an argument for the epistemic acceptability of giving up logical laws. His case against analyticity might be said to constitute just such an argument. But we have already seen that that argument does not succeed. Moreover, Gibson (1988, 42) rejects the claim that the argument against analyticity is used to support holism. As indeed he must, since, on Gibson's interpretation, Quine uses holism to argue against analyticity,[62] using the argument against analyticity to support the argument for holism would render both cases viciously circular.

Things are considerably worse for Quinean holism, however. According to the holist, the only reason for believing anything to be true is the predictive success gained for empirical science by including that statement in our web of beliefs. One consequence of this is that we can never have any reason to believe any sentence of mathematics which does not, in some way, help us deal with experience. This means that we have no reason for believing true those portions of mathematics which are unapplied.[63]

Quine himself seems at ease with this consequence of his holism. He has referred to the unapplied portions of mathematics as "uninterpreted systems,"[64] and dismissed unapplied mathematics as "mathematical recreation without ontological rights."[65] But this, surely, is a hard pill to swallow. After all, the mathematics which is unapplied is discovered by the same mathematicians using (broadly speaking) the same methods as that which is applied. Surely if the methodology and the practitioners are the

---

[62] The case, in brief, is that if all statements are revisable for empirical reasons, there is no room for a class of statements which have no empirical content, and which, therefore, would not be revisable for empirical reasons. See also Quine, 1961, 43.

[63] See also Maddy (1990, 31), and Jody Azzouni, Metaphysical Myths, Mathematical Practice: The Ontology and Epistemology of the Exact Sciences (Cambridge: Cambridge University Press, 1994),71-3 for similar arguments.

[64] See his "Review of Parson's Mathematics in Philosophy," Journal of Philosophy (1984): 788.

same in both cases, we do not have to wait upon the physical sciences to apply an as yet unapplied piece of mathematical reasoning before holding that it is true?

Nor is this the worst of it. Suppose a piece of mathematics which is currently unapplied (and hence not true) were to be needed for some explanatory purposes in some future science. It would then become true. Suppose though that a science even further in the future could do without it. It would then return to being not true. Empirical holism, then, not only strips mathematical truths of their necessity, it also seems to allow the possibility of vacillations in their having truth-values at all. This seems an unwelcome result.

Resnik attempts to avoid the problem of the lack of empirical justification of unapplied mathematics by appeal to the possibility of such parts of mathematics being required for future science.[66] The fact that such portions of mathematics may be needed by future science is what justifies us in believing them true. But why should the mere possibility of scientific application be relevant to the justification of mathematics? At least, in Quine's holism, we can see how justification is supposed to flow: the success of science justifies us in—and indeed commits us to—believing in the existence of those entities over which scientific theorizing must quantify; science ineliminably quantifies over mathematical entities, and so we are justified in believing in such entities. The justification for belief in portions of mathematics which are not now—and might never be—needed for scientific purposes cannot flow along such lines. The current success of science clearly confers no justification on such portions, and the fact

---

[65] See his "Reply to Parsons," in Hahn and Schilpp, (1986): 400

[66] Michael Resnik, Mathematics as a Science of Patterns (New York: Oxford University Press, 1997), 146-7.

that they may never be needed by science shows that we cannot count on future applications in order to justify them. Moreover, even if we could count on future science making use of currently unapplied mathematics, how would this be relevant to the current justification of such mathematics? Surely this would be relevant only to its future justification. It seems that Resnik's suggestion can, at best, allow only for unapplied mathematics being potentially justified, and this potential justification can only become actual at the time of actual application—which may never occur. This leaves the holist with the result that currently unapplied mathematics is both currently unjustified, and perhaps never to be justified. This seems hopelessly mistaken.

Michael Devitt suggests the following response for the holist.[67] Why not say that certain axioms are justified via scientific application, and the rest of mathematics is justified by deduction from these axioms? Wouldn't this get the holist out of his current troubles? The suggestion is an appealing one, but, unfortunately, it will not work. As Devitt noted when presenting this response, the attempt to capture all mathematical truths within a single system defined by a consistent set of axioms received its death-blow from Gödel. In an attempt to avoid the problem presented by the fact that there is no complete and consistent formal system, Devitt suggests that, when faced with truths not implied by our initial set of axioms, we simply change systems. Faced with a mathematical truth, p, which has no scientific application, and which our axiomatic system fails to capture, we come up with another set of axioms from which we can deduce p; p is now justified relative to this new set of axioms.

---

[67] Personal communication.

There are two problems with this move, however. The first one is that the very fact that the holist desires a derivation of p shows that he knows it is justified independent of any scientific application, or of any new set of axioms. If, as Devitt's suggestion has it, the truths of unapplied mathematics are justified by the scientific application of axioms from which they are derived, why should the holist care at all about an unapplied piece of mathematics which is not derivable from our original set of axioms? Obviously because he knows that p is already justified, independent of any scientific application the new set of axioms might have, and, further, that if the holist were not somehow to show p to be justified from within his holism, we would have a straightforward counterexample to holism. But if p is justified independent of any application, Devitt's suggestion must be mistaken in its claim about how justification in unapplied mathematics proceeds.

Secondly, there is a problem here regarding the direction in which justification flows. The holist, whatever else he holds, must hold that justification flows from beliefs of immediate experience to more abstract beliefs, less immediately connected to experience. But the jump from one axiomatic system to another suggested here reverses the direction in which justification flows. The only reason we change systems in our example is because we know that p, even though p is both unapplied and is not deducible from the initial set of axioms. In order to capture p, we pick one (or more) new axiom(s), thus defining a new axiomatic system whose sole justification is to allow the deduction of p. In other words, it is not that our new system of axioms justifies p; it is p that justifies the new set of axioms. We may indeed be able to make it appear as though p is justified by the axioms—p is, after all, deducible from them—but no set of axioms which did not imply p would have been acceptable; we have simply worked backwards, and derived our axioms from p. Justification here flows in the opposite direction to that allowable by holism; thus Devitt's attempt to revive holism fails.

There is, finally, a general worry about the coherence of the holism which underwrites this argument against the a priori. In outline, the worry is that Quine's web model, according to which all sentences are empirically revisable, must assume that at least some logical principles are immune to empirical revisability, and thus renders itself incoherent. Katz (1998, 73) identifies non-contradiction, universal revisability of sentences, and simplicity as the constitutive principles of Quine's belief-revision epistemology. Without non-contradiction, there would be no reason that recalcitrant experience would require re-distribution of truth values across the web; without universal revisability of sentences, there would remain some set of sentences immune to revisability in the face of recalcitrant experience, and Quine's thoroughgoing empiricism would be compromised; without simplicity, the class of potentially revisable sentences for any given recalcitrant experience would be too large. These principles must therefore be part of the argument for the revision of any given belief in the face of recalcitrant experience. Katz (1998, 73) then presents what he calls the "revisability paradox." Since the constitutive principles are part of the argument for the revision of any given belief, they cannot themselves be revisable:

> Any argument for changing the truth value of one of the constitutive principles must have a conclusion that contradicts a premise of the argument, and hence must be an unsound argument for revising the constitutive principle.

Thus, either empirical holism contradicts itself by holding its constitutive principles immune from revision, or else it licenses contradictory arguments in favor of revising these principles.[68]

So, not only does the argument from holism beg the question against the believer in a priori knowledge, holism itself leads to some fairly undesirable conclusions. There are two conclusions to be drawn. The first is that the argument from holism against a priori knowledge does not succeed. The second is that (on the assumption that no brand of empirical Platonism which requires that we directly perceive mathematical objects is defensible) apriorism is the only tenable epistemology for the Platonist.

### 2.3.1.

### The Argument from Naturalized Epistemology (Bonjour's Interpretation)[69]

Both Bonjour (1998) and Gibson (1988) see a case against the a priori in Quine's naturalized epistemology. In his recent In Defense of Pure Reason, Lawrence Bonjour undertakes a defense of a priori knowledge in the face of Quine's attacks on that notion. In the course of this defense, Bonjour examines an argument which moves from a naturalized epistemology to a rejection of the a priori. Bonjour (1998, 82) identifies the case against the a priori stemming from a naturalized epistemology as follows: once we have settled for a naturalized epistemology, there is "no reputable cognitive endeavor

---

[68] Richard Warner presents a similar argument to the effect that unless logic is assumed to be a priori, it becomes impossible to say whether an experience does or does not require revision of some part of the web. See his "Why Is Logic A Priori?" The Monist, 72 (1989): 46-7.

[69] In altered form, section 2.3.1 appears as "Naturalized Epistemology, Normativity, and the Argument Against the A Priori," in the internet journal Essays in Philosophy, 2002, <www.humboldt.edu/~essays>.

that requires any sort of a priori appeal." Showing that this is so would presumably involve showing that epistemology (or, at least, whatever part of it that can be shown to be reputable) can be done without such an appeal. This, of course, is the task Quine sets himself in "Epistemology Naturalized."[70] The central claim there is that we should stop dreaming of a first philosophy, one which provides a foundation for science, and settle instead for an account which explains how we in fact arrive at our theories of the world, given only the meager sensory input. We should settle, that is, for descriptive psychology rather than epistemology as traditionally conceived. This new enterprise will be thoroughly empiricist, and if it can in fact provide a satisfactory link between evidence and theory, it will have rendered otiose any appeal to a priori knowledge.

This, then, is the argument against a priori knowledge, from naturalized epistemology. The main business of the present section is to argue that it fails. Unlike Bonjour and Kim, [71] however, I will not contend that a Quinean naturalized epistemology leaves out normative concepts essential to epistemology, and thus is not properly an epistemology. To the contrary, I shall attempt to show that there is no quick rejection of the argument from naturalized epistemology along such lines. To this end, I shall first show how a Quinean naturalized epistemology makes room for the normative notions of truth and justification. Following this, I shall argue that, even allowing that Quine's epistemology can accommodate the normative, the argument from a naturalized epistemology does not succeed.

---

[70] Quine, "Epistemology Naturalized," in Ontological Relativity and other Essays (New York; Columbia University Press, 1969).

## 2.3.1.1

## Naturalized Epistemology and Justification

Kim's (1993) is perhaps the best-known criticism of Quine's project. Psychology, according to Kim, is a descriptive science. Epistemology, on the other hand, is fundamentally a normative enterprise. Quine, Kim (1993, 333) contends, is asking us to give up the entire framework of justification-centered epistemology, and to put in its place a purely descriptive, causal-nomological science of human cognition. Kim responds that due to the fact that our concept of knowledge is inseparably tied to that of justification, if justification drops out of epistemology, so too does knowledge. For the epistemologist to drop both justification and knowledge is for him to go out of business.

Kim considers a possible line of defense for the naturalized epistemologist along the following lines. Contrary to first appearances, justification does not drop out of a naturalized epistemology. This is because Quine's conception of the role of the new epistemologist has him investigating how evidence relates to theory, and this relationship is a justificatory one. Thus, the response concludes, the naturalized epistemologist has his own brand of justification, and the objection is misplaced. However, Kim argues that the only relation between sensory input and theoretical output to which the naturalized epistemologist is entitled is a causal one, and not an evidential one. The causal relation cannot be an evidential one because the causal

[71] Jaegwon Kim, "What is 'Epistemology Naturalized'?" Reprinted in The Theory of Knowledge: Classical and Contemporary Readings, edited by Louis Pojman (Belmont, Ca.: Wadsworth, 1993). All page references to Kim's article refer to this reprint.

mechanisms connecting input and output will vary across species. A truly evidential relation abstracts from such species-dependent features, concerning itself only with the degree to which evidence supports theory—and this is a relation holding between contents, rather than between causes and effects. But, since this is exactly what the naturalized epistemologist cannot have, he cannot capture justification, and the initial objection stands.

The claim, then, is that in becoming part of natural science, epistemology loses all claim to normativity, and thus to justification. But how true is this? Certainly Quine himself has denied jettisoning the normative. In his "Comment on Lauener,"[72] he writes that "The normative is naturalized, not dropped." Normative epistemology, we are told a few lines later simply is scientific method. In the Hahn-Schilpp volume, Quine, in reply to White, writes:

> Naturalized epistemology does not jettison the normative and settle for the indiscriminate description of ongoing processes. For me normative epistemology is a branch of engineering. It is the technology of truth-seeking, or, in more cautious epistemic terms, prediction.[73]

But just how is descriptive science normative? One possible answer to this question makes appeal to natural selection.[74] Since natural selection selects belief-forming processes that generally lead to true beliefs, so this answer goes, the beliefs we actually

---

[72] Quine, "Comment on Lauener," in Barrett and Gibson, 1990, 229.

[73] Quine, "Reply to White," in Hahn and Schilpp, 1986, 664-5.

[74] See, for example, Hilary Kornblith, "What is Naturalistic Epistemology?" in Naturalizing Epistemology, ed. Hilary Kornblith (Cambridge, MA.:MIT, 1987).

have are the ones we ought to have. The appeal to natural selection thus captures the normative. Further, since the theory of evolution is part of our naturalized scientific picture of the world, appeal to it is sanctioned by a Quinean naturalized epistemology. So, it would appear, there is no problem for a naturalized epistemology.

The above response, however, relies on a tendentious account of what natural selection does. Natural selection selects belief-forming processes which have survival value. Such processes need not coincide with processes which generally lead to true beliefs. A more cost-effective method of enhancing prospects of survival might involve selecting for a fairly high number of mistaken beliefs (for example, that any moving object over a certain size is likely to be a predator), rather than selecting for all or mostly true beliefs. Granted this view of natural selection, the claim that it ensures that the beliefs we have are the ones we ought to have is a dubious one.

A further problem with the suggestion that the normative can be captured by appeal to natural selection arises from consideration of highly theoretical beliefs. What survival value do our beliefs about trans-finite sets have, for example? If, as seems likely, such beliefs have little or no survival value, how do we account for their normativity on the evolutionary picture? It would appear that the naturalizer of epistemology will have to look elsewhere in his attempt to capture the normative.[75]

Quine's own answer to the problem of normativity, however, runs in a different direction to the evolutionary one. We find the beginnings of his answer in the reply to White:

---

[75] The fact that we are capable of having such beliefs can, perhaps, be explained by means of an exaptation, but the normativity of such beliefs would remain unexplained by such a move.

> [Naturalized Epistemology] draws upon psychology in exposing perceptual illusions, and upon cognitive psychology in scouting wishful thinking. It draws upon neurology and physics, in a general way, in discounting testimony from occult or parapsychological sources (in Hahn and Schilpp, 1986, 664-5).

Similarly, in the "Comment on Lauener," natural science is described as "conspicuously normative" in its counsel to mistrust soothsayers and telepathists (Quine, 1990a, 229). So, natural science is normative, at least to the extent that it rules out certain avenues which might be thought to help us in our dealings with the physical world. But there is more to the normativity of science than just this. As we ought to expect from a pragmatist philosopher like Quine, the normativity of natural science is a matter of how successful its predictions turn out. Science, Gibson (1988,75) tells us, "is justified in Quine's eyes by its measuring up to observation and prediction." In other words, science, or any part thereof, is justified to the extent that it successfully enables us to control and/or predict the workings of the physical world. The point is underscored in Quine's "Comment on Quinton."[76] Writing on the status of the law of non-contradiction, Quine explains that the reason we do not, in practice, revise this law is that "without it we would be left making mutually contrary predictions indiscriminately, thus scoring a poor ratio of successes over failures" (Quine, 1990b, 309) .

We have then, the outline of a Quinean response to Kim's objection. Empirical science indeed has room for normative concepts such as justification. Sentences, or theories, are justified to the extent that their part in our overall science facilitates successful observation and prediction. Our beliefs are justified to the extent that they

[76] Quine, "Comment on Quinton," in Barrett and Gibson, (1990), 309.

enable us to deal with, or theorize about, the world. The claims of clairvoyants are unjustified because their ratio of success over failure is poor. The case is just the opposite with quantum physics, or with the law of non-contradiction.

Moreover, for Quine, there is no further question about how scientific principles and methodology are justified. Once the goal of a first philosophy has been abandoned natural science is "not answerable to any supra-scientific tribunal, and not in need of any justification beyond observation and the hypothetico-deductive method."[77] Quine explains why it is not:

> The naturalistic philosopher begins his reasoning within the inherited world theory as a going concern. He tentatively believes all of it, but believes also that some unidentified portions are wrong (Quine, 1981c, 72 emphasis added).[78]

Since, on Quine's conception, philosophy is a part of the inherited world theory, it cannot assume the role of a "supra-scientific tribunal." Science, then, is a self-justifying process, which neither has nor needs the a priori justification of a first philosophy. It would appear, then, that Kim's criticism is off the mark. Quine's naturalized epistemology can indeed make room for the normative conception of justification.

One might respond in Kim's defense along the following lines. The foregoing accommodation of justification within naturalized epistemology works by equating the justification for a claim (or body of claims) with the success of the predictions which can be made on the basis of that claim (or body of claims). However, it is, as the history

---

[77] Quine, Theories and Things, (Cambridge, MA.: Belknap/Harvard, 1981), 72.

[78] For further explicit statement of how Quine's epistemology assumes in advance the ontology of science, see his "Empirical Content," in Quine 1981: 24.

of science shows, possible to obtain a high degree of predictive accuracy from a body of statements which are in fact false. This being so, justification conceived of as high predictive success fails to be truth-conducive. But the main reason for caring about justification is to try to order our beliefs in such a way as to have as many true beliefs, and as few false ones, as possible. So, if predictive accuracy is not truth conducive, then it would appear that the Quinean ersatz justification does not give us what we want from justification. It thus remains to be shown that a naturalized epistemology can make room for the normative.

The response, however, fails. In the first place, the relationship between predictive accuracy and truth is, in the relevant respect, identical to the one between justification and truth. Just as we can come to believe a body of false claims to be true because they yield accurate predictions, so we can come to believe false statements to be true because, in our imperfect knowledge situation, the beliefs we already have (apparently) justify them. In both cases it is possible to have a body of true beliefs (correct predictions in the one, and the true beliefs apparently justifying the false claim in the other) which offer reason to think that another (false) statement is true. Thus, if the above response sufficed to show that predictive accuracy is not truth-conducive, a similar move would show that neither is justification.

In the second place, the claim that predictive accuracy is not truth-conducive is mistaken. Although false theory can, in the short run, yield accurate predictions, the probability of a hypothesis increases with its predictive success. Since the probability of a hypothesis is the probability of that hypothesis being true, predictive accuracy is, over the long run, truth conducive.[79] The response made on behalf of Kim thus fails, and there seems little wrong with Quine's claim to be able to capture the normativity of justification in a naturalized epistemology.

### 2.3.1.2

### Naturalized Epistemology and Truth

The second criticism of the argument Bonjour attributes to Quine is due to Bonjour himself. He (1998) presents an argument based on the nature of the relationship between truth and justification. Psychology, he writes (1998, 85), concerns itself with causal relations. It has no place for the concept of a reason for thinking some claim is true. A few pages later (1998, 87), in the context of a discussion of Quine's ability to deal with skepticism, Bonjour describes what he takes to be one of the fundamental questions of epistemology: Is the justification that is available for the belief in question genuinely adequate to show that it is (at least) likely to be true? Since, in Quine's naturalized epistemology, justification is a matter of how a belief is situated in the web, Bonjour sees Quine as having nothing whatever to say in answer to this question. Naturalized epistemology, then, fails to address a fundamental question of epistemology, and thus is inadequate as epistemology. Although Bonjour does not elaborate, the point seems to rest on the assumption that how a given belief is situated

[79] Thanks to Michael Levin for pointing out this response to me.

with respect to other beliefs is merely a matter of individual psychology, and hence has no bearing at all on whether the belief is likely to be true. [80]

But Bonjour is mistaken to think that the position of a given belief within the Quinean web of belief has nothing to do with the probability of that belief being true. For Quine, the only good reason we can have for including a given belief in our web of beliefs is its passing the hypothetico-deductive test of science. And, passing this test increases the probability of that belief being true. Thus, the position of a given belief, properly included in the web of belief, indeed does offer a reason for believing that that belief is (at least likely to be) true, and the objection is misplaced.

I have thus far argued that, contra Kim and Bonjour, there is room within Quine's naturalized epistemology for such normative concepts as truth and justification. The argument from naturalized epistemology, then, cannot be dismissed on the grounds that the epistemology which underwrites it is essentially misguided. But can Quine's naturalized epistemology fund an argument against the a priori?

### 2.3.1.3

### Quine's Conception of the A priori in "Epistemology Naturalized"

The argument from naturalized epistemology is simply that such an epistemology can give you all you could fairly ask from epistemology, and do so without leaving any reputable work for the a priori. The question, then, is what reasons does the Quine of

[80] Kim (1993, 329) also sees the question of what conditions a belief must meet if we are justified in accepting it as true as one of the two fundamental questions of epistemology. However, unlike Bonjour, he does not go on to make much of this in connection with Quine's naturalized epistemology.

"Epistemology Naturalized" offer for thinking that there is no reputable work for the a priori to do?

One possible answer to this question is that naturalized epistemology by definition rules out the a priori, and that to accept such an epistemology is ipso facto to reject the a priori. This, however, will not do. One can of course, simply refuse, without argument, to accept any notion of a priori knowledge into one's epistemology. However, if one takes this road, one merely asserts one's aesthetic preferences. A "victory" so gained over the apriorist is a hollow one indeed. Moreover, this is not the route taken by Quine.

Consider the following, which occurs in a discussion of what Quine sees as the goals of traditional epistemology: "[t]he old tendency was due to the drive to base science on something firmer and prior in the subject's experience; but we dropped that project" (Quine, 1969, 87). Consider also the following remarks, from a passage where Quine (1969, 75-76) is defending the methods of naturalized epistemology from the charge of circularity:

> If the epistemologist's goal is validation of the grounds of empirical science, he defeats his purpose by using psychology or other empirical science in the validation. However such scruples have little point once we have stopped dreaming of deducing science from observations.

Traditional epistemology, with its a priori methodology, Quine contends, had the goal of providing a priori and firm (i.e., certain) foundations for science. There are good inductive reasons for thinking that this cannot be done. Thus, Quine's suggestion that we give up the traditional epistemological project, so conceived, and settle for psychology.

But this is too quick. If one's understanding of the a priori requires, along Cartesian or Carnapian lines, that it provide foundational, and certain knowledge, and one is also persuaded by the failure to provide any such foundations, one might well think that in providing a naturalized epistemology, one which is conspicuously normative, one had thereby provided an argument against the a priori. But to show that a naturalized epistemology can make room for the normative is not to show that there is no reputable cognitive endeavor left for the a priori, and this for two reasons.

Firstly, In showing that a naturalized epistemology can lead to beliefs which are likely to be true and justified, one does not thereby show that passing the hypothetico-deductive test of the physical sciences is the only method which can lead to such beliefs. The possibility remains that there are, in addition to the test of physical science, a priori methods of arriving at true and justified beliefs. Secondly, the apriorist does not necessarily concern himself solely, or even at all, with the foundational project described by Quine. If there are other, more attainable goals for the apriorist, then to reject traditional epistemology on the grounds that it cannot offer firm foundations for science is to throw the baby out with the bathwater. Even if we cannot have any a priori knowledge which is foundational in this sense, this has no tendency to show that there can be no a priori knowledge whatever. And, as I shall contend, when one restricts oneself to a more moderate conception of the a priori, there are more attainable goals.

According to the conception of the a priori on which Quine set his sights in "Epistemology Naturalized," a priori knowledge is thought to lead, infallibly, to certainty. On a more moderate conception, a priori knowledge is simply knowledge which does not depend on experience. Thus, for there to be a priori knowledge on this conception, it is enough if there are processes which yield beliefs without the warrant for those beliefs depending on any particular characteristic of experience.[81] The most conspicuous example of such knowledge is mathematical knowledge. For sure, some experience is required, in order to acquire the requisite concepts, but beyond concept acquisition, the warranting of mathematical beliefs requires no particular experience. It requires only intellectual understanding. Although one may encounter a mathematical proof by reading it from a chalkboard, or by seeing it in a textbook, or by hearing it uttered, these causal origins of the belief are not justification for believing it. One comes to have justification for believing in the truth of the conclusion not through any such experience, but through understanding the proof, and thereby coming to understand that the conclusion must be true. This kind of intellectual understanding is not dependent upon any particular characteristic of anything properly called experience, and thus, on a moderate conception of the a priori, mathematical knowledge is knowledge a priori.

Of course, Quine would disagree with this characterization of mathematical knowledge, as he sees it rather as part of our empirically justified web of belief, no different in kind from the paradigmatically a posteriori knowledge yielded by physics. But we have seen that holism is both unworkable as an epistemology, and unable to fund the conclusion that the only knowledge there is is knowledge a posteriori. In any

[81] I take this characterization of the a priori from Paul Boghossian and Christopher Peacocke's introductory essay to their (2000): 1-2.

event, my contention here is just that there is no sound argument from naturalized epistemology against the a priori, not that Quine does not have other reasons for thinking as he does. It suffices, in order to show this, to show that Quine is working with an unnecessarily strong notion of the a priori, and that nothing he says in "Epistemology Naturalized" rules out a more moderate conception, one which does not tie its fate to the attempt to provide absolutely certain, a priori, foundations for the physical sciences. And this is done by offering both the moderate account of the a priori sketched above, and the prima facie plausible claim that mathematical knowledge is a priori on this conception.

There is, however, a suggestion of Gibson's, which, if sound, would provide a way of resurrecting the argument from naturalized epistemology. The suggestion is that traditional epistemology is in fact incoherent, because epistemology presupposes ontology. Epistemology presupposes ontology since it presupposes the existence of both the external world and physical nerve endings which yield information about it (Gibson 1988, 66; Quine 1981b, 24; 1981c, 72; Quine 1969, 82-3). Building on this, the case against the a priori would run something like the following.

1. Traditional, a priori, epistemology seeks knowledge which makes no ontological assumptions.
2. In fact, epistemology presupposes ontology.
3. Thus traditional, a priori, epistemology is incoherent.
4. Therefore, naturalized epistemology is the only coherent epistemology we have.

5. Naturalized epistemology rejects the possibility of a priori knowledge.
6. Therefore, our only coherent epistemology rejects the possibility of a priori knowledge.

But as an argument against the a priori, this clearly cannot work. For not all epistemology presupposes the existence of the external world. Quine's naturalized epistemology does, it is true, but the fact that naturalized epistemology presupposes the ontology of natural science is not to say that epistemology presupposes ontology. Reliabilism, for example, does not presuppose ontology. To see this, consider the reliabilist's answer to the skeptic. That answer is simply that if our beliefs are true and reliably caused, then we have knowledge of the external world. The reliabilist does not go the further step and claim that we know that these beliefs are true and reliably caused. He does not hold that we know that we know that there is an external world. He thus does not hold that we know that we know that we have nerve endings or that there is an external world. That is, his epistemology does not make ontological assumptions.

Further, to assume the truth of (2) is to assume the truth of naturalized epistemology in an argument against the a priori. This of course, would beg the question against the apriorist. For both these reasons, then, premiss (2) has not been established; it thus cannot be used to support either premiss (3) or (4), or the conclusion, (6).

Moreover, premiss (1) is false too. Much realist philosophy of mathematics proceeds on the assumptions that the objects of mathematical knowledge exist, and do so abstractly. Making such assumptions is quite consistent with the realist's claim to be discovering truths a priori. Of course, the realist might go on to give arguments for his assumptions. Doing this, however, is a separate exercise, and it does not at all undermine the claim that one can discover truths about abstract objects while only assuming, as opposed to establishing, that they exist. Such knowledge would be conditional, but it would be knowledge. Contra Gibson, then, traditional a priori epistemology is quite consistent with the presence of ontological assumptions. Thus, the attempt to resurrect the argument from naturalized epistemology on Gibson's suggestion fails. The conclusion must be that, so far at least, no legitimate case has been made against the a priori from a naturalized epistemology.

### 2.3.2

### The Argument from Naturalized Epistemology (Gibson's Interpretation)

For Gibson, like Bonjour, the case against the a priori arising from Naturalized epistemology rests upon the ability of a naturalized epistemology to successfully take the place of traditional epistemology. If then, Quine's "scientistic turn" (Gibson, 1988, 29) in epistemology can be justified, this would render otiose both traditional epistemology, and its appeal to a priori methodology. According to Gibson's (1988, 29) interpretation, there are two parts to the justification for the scientistic turn:

> First, he [Quine] argues that skepticism about science presupposes science. Second, he argues that science needs no justification beyond measuring up to the demands of observation and hypothetico-deductive method.

However, if these arguments indeed jointly comprise the intended justification for a naturalized epistemology, then the latter is in trouble, since neither of the supporting arguments are themselves well-supported.

Consider the second part of the justification. The Quine of "Epistemology Naturalized" does indeed urge the claim that science needs no justification beyond measuring up to the demands of observation and hypothetico-deductive method. But when one recalls that by 'science', Quine means not merely the physical sciences, but our entire body of knowledge, the claim becomes extremely problematic. The claim that (e.g.) mathematics and logic need no justification other than their measuring up to the demands of observation and the hypothetico-deductive method is extremely implausible. The only way in which the Quinean might support such a claim is by appeal to his holism. But we have already examined this, and found it unacceptable. There is, then, no support for the claim that science, read in Quine's preferred all-encompassing way, needs no further justification than that mentioned.

If, on the other hand, we understand the claim to pertain only to the physical sciences, rather than to our entire body of knowledge, then it becomes hard to see how the claim, even if true, could support a naturalized epistemology. Quine's brand of naturalized epistemology, after all, is supposed to render appeal to a priori justification otiose. If the argument in its favor shows, at most, that there is no need for an a priori justification of the physical sciences, its job is only half done. The defender of the a priori need only gesture to the formal sciences, and observe that, for all that has been said, they are still in need of a priori justification, and that the attempt to render a priori justification otiose has thus failed. On either interpretation of 'science', then, the attempt to support a naturalized epistemology via the claim that science needs no justification other than that provided by observation and the hypothetico-deductive method is seen to fail.

Nor does the first part of the justification fare any better. That had it that the justification for naturalized epistemology rests on the further claim that skeptical doubts about science presuppose science. Note that in what follows, I am not arguing that Quine's position itself leads to skepticism,[82] only that he has not shown that the skeptic presupposes science.[83] What is Quine's argument that the skeptic presupposes science? Quine (1981a, 22), urges that skeptical doubts "would still be immanent and of a piece with the scientific endeavor." So, since skeptical doubts must arise within science, they must share at least part of a common ontology. The skeptic can agree that he assumes a scientific ontology, but will add that he does this only in order to refute it. Since Quine agrees that the skeptic is quite within his rights to do this,[84] how is the skeptic's assumption supposed to be in any way illegitimate, or harmful to his case? Gibson's (1988, 60) answer is that to suppose that all of our ontological commitments could be simultaneously doubted involves a position which is "a transcendental one, and thus an incoherent one."

[82] For arguments to the effect that Quine's position does lead to skepticism, see Katz (1990, 298-307), Bonjour (1998, 89-97), and Barry Stroud, The Significance of Philosophical Skepticism; Oxford: Oxford University Press, 1984), 209-54.

[83] To be more precise, since the skeptic assumes science in order to refute it, and since Quine agrees that he is within his rights to do so, I am going to argue that the Quine has not shown that the skeptic assumes science in any illegitimate way.

[84] Quine "The Nature of Natural Knowledge," in Mind and Language, ed. S. Guttenplan (Oxford: Clarendon, 1975), 68.

But how are transcendental doubts supposed to be incoherent? Quine's naturalized epistemology, as we have seen, assumes an ontology of nerve endings, and physical forces which impinge upon these. From within such an epistemology, it is certainly true that one cannot coherently doubt all of one's ontological commitments simultaneously. For even arguing from physical illusions requires some way of telling illusions from veridical perception, and this, in turn, requires appeal to nerve endings and impinging forces. But to say that transcendental doubts are incoherent when they arise within an epistemology which rules out such doubts is not to say that transcendental doubts are incoherent simpliciter. This would only be so if Quine had shown that the entire traditional epistemological project had been misconceived. If Quine had shown that, then the skeptic would be in the position where the only coherent epistemology available was one in which his doubts became incoherent. But the furthest that Quine goes toward arguing that traditional epistemology is misconceived is to argue against the foundationalist project as conceived by Descartes and by Carnap. As this does not suffice to show that the entire project of traditional epistemology is misconceived, it does not show that transcendental doubts are incoherent. Without it being established that the only epistemological game in town assumes the ontology of science, Quine's response to the skeptic begs the question against him. It cannot, therefore, be used to ward off the skeptical reductio of science. And, since it cannot, the justification which Gibson offers for the scientistic turn in epistemology (depending, as it does on the claim that skeptical doubts presuppose science) does not suffice.

There are two other readings of Quine's claim that our doubts "would still be immanent and of a piece with the scientific endeavor" but neither are threatening, either to skepticism or to apriorism. If we follow Quine in seeing 'science' as covering all of our intellectual endeavor, then it is obviously true that any part of the intellectual endeavor must arise within that same intellectual endeavor. But it is less than clear how this is supposed to show the impossibility of a first philosophy, or of a priori skepticism, or of an a priori refutation of the skeptic. One who raises the possibility of our entire network of experiences being merely a dream seems to have doubted all of his ontological commitments simultaneously. If this is all that is required for transcendental doubt, then nothing that Quine has said shows the incoherence of such doubt, nor the incoherence of the possibility of an a priori resolution of it.

If, on the other hand, transcendental skepticism requires somehow doubting from a position which does not arise within any intellectual framework, then of course this is, if not incoherent, at least problematic. But why this result might be supposed to be of any significance in the debate over the a priori would be quite mysterious.

Gibson (1988, 31), however, sees another argument in Quine's writings which, if sound, show that transcendental doubts are incoherent:

> Holism ushers out first philosophy: theoretical terms can neither be defined nor translated into terms of immediate experience, and neither the analytic nor the a priori (in general) can be isolated in any absolute manner—any sentence can be held true come what may. Unregenerate realism ushers in naturalized epistemology: skeptical doubts are scientific doubts [emphasis in original].

The argument seems to be that, since the a priori has already been ruled out by the arrival of holism, and since "unregenerate realism" has ushered in naturalized epistemology, there is nowhere outside of science from where the skeptic can doubt.

Since this is so, his doubts arise from within, and are thus part of, the scientific endeavor.

But we have already seen that the arguments against analyticity and the a priori fail to show that these cannot be "isolated in any absolute manner," so there is no reason to think the skeptic might not help his case by having recourse to a priori doubts. Moreover, given that Quine defines "unregenerate realism" as "the robust state of mind of the natural scientist who has never felt any qualms beyond the negotiable uncertainties internal to science,"[85] it is hard to see how unregenerate realism could legitimately usher in naturalized epistemology. Since naturalized epistemology comes as a package with ontological commitments, ushering this in is thereby to usher out skepticism. But since the attitude which ushers out skepticism is simply one that has never taken skeptical doubts seriously, the skeptic has cause to complain of unfair treatment. The point here is that if skepticism has not legitimately been ushered out, then an epistemology which assumes what the skeptic calls into question, i.e., a naturalized epistemology, cannot have been ushered in.

According to Gibson's interpretation of the argument from naturalized epistemology, that argument was itself supported by the further claims that skepticism about science presupposes science, and that the claims of science need no justification beyond measuring up to the demands of observation and hypothetico-deductive method. We have seen that these latter claims themselves stand in need of further support, and thus cannot be used to support an argument against the a priori.

[85] "Five Milestones of Empiricism," in Quine (1981), 72.

Granted this conclusion, together with the conclusions of previous sections, then, we must conclude that Quine has offered no compelling case against the moderate construal of the a priori outlined here. The next step is to argue that mathematical knowledge is itself a priori. To this end, Chapter Three will argue defend the traditional conception of mathematical proof as yielding knowledge a priori, and Chapters Four and Five will articulate and defend a reliabilist account of a priori mathematical intuition.

# Chapter Three

## The Apriority of Non-Basic Mathematical Processes

### Introduction

My goal in this chapter is, once more, purely defensive. I have argued that there are no compelling reasons offered by Quine for his rejection of the a priori. To make a cogent case for a reliabilist account of mathematical apriorism, two further steps are required. Chapters Four and Five will address the question of whether the process of mathematical intuition can reliably yield a priori knowledge of basic mathematical truth. For the present chapter, we will be concerned with non-basic (i.e., inferential) mathematical knowledge, usually obtained via the process of mathematical proof. It is sometimes assumed[86] that if a satisfactory account of how knowledge is generated via a priori intuition can be given, the rest would be easy for the apriorist. After all, if basic mathematical knowledge is a priori, and accepted rules of inference are a priori, the game would appear to be over. However, there are important arguments against mathematical apriorism which concentrate on mathematical proof, and argue that mathematical proof does not yield a priori knowledge. The business of this chapter is to consider the attempts of Philip Kitcher and Thomas Tymoczko to argue in this way. The

[86] See, for example, Field (1989, 231).

chapter heading refers to non-basic mathematical processes, rather than simply to proofs, since fully answering Kitcher's case will involve some discussion of non-basic mathematical processes other than proof.

Apart from their focus on non-basic mathematical processes, rather than on intuition, there are other reasons for considering the arguments of these two philosophers. Firstly, the arguments of both philosophers have been extremely influential. Secondly, they both find elements of experiential warrant in the allegedly a priori field of mathematics. Part of Kitcher's case that this is so stems from the social context of mathematical knowledge, and part from the limits of human cognition. Tymoczko's arguments arise from consideration of the growing use of computer technology in mathematics, as a means of extending those limits. Thirdly, and perhaps most importantly, my answers to the objections raised by these philosophers have, in important ways, shaped my conception of what a priori knowledge must be, if it is to avoid these, and other objections. Finally, in Kitcher's case, the objections to mathematical apriorism stem from a (broadly) reliabilist epistemology. Kitcher sees his account of a priori knowledge as being the only fully developed such account, and as ruling mathematical knowledge to be beyond its scope. I address his arguments in order to show that a reliabilist epistemology does not, in fact, have this consequence.

Although Kitcher presents arguments which specifically address mathematical proof, and others which specifically address the apriorist Platonist process of mathematical intuition, he also presents some arguments that are not specific to any particular—allegedly a priori—process. These arguments instead seek to show that mathematical knowledge—regardless of the particular producing process—is not a priori. I shall deal with these arguments in this chapter, since leaving them to the chapter on intuition would leave the case for the apriority of proof awaiting the case for the apriority of intuition.

Before discussing Kitcher's analysis of the a priori, and his arguments against mathematical apriorism, I should say a little about the version of mathematical apriorism here to be defended. That position is one which is both fallibilistic and reliabilist with respect to warrant, and Platonist with respect to ontology. All aspects of the position will be used to defuse the arguments discussed in this chapter.

I begin by assuming that, until given reason to think otherwise, it is initially plausible to see mathematical knowledge as knowledge a priori. Recall, from the previous chapter, our moderate (and vague) account of a priori knowledge as knowledge the warrant of which does not depend on any particular characteristic of experience. According to this moderate conception, for there to be a priori knowledge, it is enough if there are processes which yield beliefs without the warrant for those beliefs depending on experience. And, although in coming to know something via a proof, we typically perceive that proof, it is not the perception itself that yields the warrant. Perception yields the access to that warrant, but the warrant itself is provided by the proof. Granted that this is at least initially plausible, if it is possible to defuse those reasons offered for denying that mathematical knowledge is a priori knowledge, we can justly maintain the position that it is a priori. In defusing those reasons, we will also specify more clearly what is meant by saying that proof is an a priori process. Thus will the charge of vagueness be avoided: although our definition of a priori warrant will not itself be sharpened, in saying how proof provides warrant, independently of experience, we thereby provide an account of what it means for the process of proof to be properly said to be a priori in this sense, which is not itself vague.

The process of following a proof, then, will properly be counted as a priori if it does not require experience for its warrant. If it is also repeatable and reliable, it will, by reliabilist lights, count as an a priori warranting process. It is not required that one have access to one's warrant.[87] It is enough that the process be reliable, whether we know it to be so or not.

The Platonist aspect of the position has it that a proof is an abstract, unchanging structure, the premisses of which necessarily imply the conclusion. On such an ontological picture then, correctly following a proof gives a completely reliable warrant for belief in the conclusion (though, as we shall see, there may well be questions about whether we have correctly followed a given proof). Kitcher does not share these ontological assumptions. We cannot, then, use these assumptions to attack Kitcher's preferred ontology of mathematics. Since, however, he is arguing that there is no coherent brand of mathematical apriorism, we are quite within our rights to make use of arguments flowing from a Platonist ontology in order to respond to his attack.

With respect to the fallibilist aspect of the position, things are a little more complicated. As will emerge, I believe that the combination of Platonism and reliabilism results in the conclusion that genuinely a priori mathematical processes cannot change their status. That is, they cannot move from being warranting processes

[87] Kitcher (2000, 66) agrees that the subject is not required to be able to cite some set of justifying propositions in order to be warranted. However, he also holds (1983, 20) that the mere correct functioning of a reliable process will not be enough to warrant a subject's belief that p, if for example, he believes (incorrectly) that the producing process is unreliable. As should be evident from my discussion of the Bonjour cases, I do not follow Kitcher in this.

to being non-warranting processes. So, the fallibilism in question is not one that holds that a priori mathematical processes are fallible (in the sense that they may fail to confer warrant). Two things are meant by my fallibilism about the a priori. The first is that warranting a priori processes can produce beliefs which are false, but warranted. Thus, a priori warranted beliefs are fallibly warranted, and the warrant attaching to them is defeasible. However, since I hold that what is warranted a priori is not empirically defeasible, I thereby hold that a priori warranted beliefs can have their warrant defeated only by other a priori considerations. More on this, ever so much more on this, anon. The second aspect to my fallibilism is that, as humans, we are fallible in our ability to demarcate which processes are in fact warranting. We can believe of a process which actually is warranting that it is not, and vice versa.

## 3.1

## Kitcher's Analysis of the A Priori

Philip Kitcher (1983, 24) presents the following analysis of a priori knowledge.

> (1) X knows a priori that p iff a knows that p, and X's belief that p was produced by a process which is an a priori warrant for p.

This raises the question of what is meant by an a priori warrant. Kitcher's (1983, 24) answer to this is that:

(2) W is an a priori warrant for X's belief that p iff W is a process such that, given any life e, sufficient for X for p,

(a) Some process of the same type could produce in X a belief that p.

(b) If a process of the same type were to produce in X a belief that p, then it would warrant X in believing that p.

(c) If a process of the same type were to produce in X a belief that p, then p.

Warrant itself is left unanalyzed. This is deliberate. Kitcher (1983, 18) sees (2) as a condition which must be satisfied by beliefs which are warranted a priori, in addition to whatever conditions must be satisfied for a belief to be warranted simpliciter. Kitcher's analysis of a priori warrant, then, is held to be independent of any particular account of whatever the third condition on knowledge turns out to be.[88] The reference to a life sufficient for X for p functions so as to acknowledge the Kantian point that experience is necessary in order to provide a believer with the relevant concepts. If one had never had experience of any mathematical concepts, one could not form any mathematical beliefs. This point on its own, however, does nothing to impugn the status of mathematics as an a priori discipline. For the relevant question to ask when considering whether a proposition can be known a priori is whether one's warrant for belief in that proposition must make essential reference to experience. If not, then the proposition is known a priori.

[88] That is, except for the fact that Kitcher (184, 15) holds that our account of knowledge must be psychologistic. His reason for this is to ensure that we do not consider someone's belief that q to be warranted merely because he believes that p and that p implies that q. These latter two beliefs must be, in some way, causally relevant to his believing that q.

As I shall be in disagreement with Kitcher both about which of (a) through (c) are properly considered necessary conditions on a priori warrant, and about whether the conditions I deem acceptable have the consequences claimed by Kitcher, it is important to be clear on why Kitcher chooses these conditions. Using Kant's account of pure intuition to develop his own account, Kitcher (1983, 23) writes:

> The same type of process must be available independently of experience. It must produce warranted belief independently of experience. And it must produce true belief independently of experience [emphasis in original].

Thus, clause (a) requires that processes which confer a priori warrant should be repeatable. Clause (b) imposes a reliability condition on an a priori warrant. If a process in fact confers an a priori warrant on a subject's belief that p, then a process of that same type would confer an a priori warrant for belief that p—regardless of the subject's experiential setting. Clause (c) is a truth condition. If some process confers a priori warrant on the belief that p, then p.

Having spelled out what he takes a priori knowledge to be, Kitcher goes on to argue that mathematical proof does not satisfy the conditions on it. One could, of course, simply refuse to consider Kitcher's arguments, on the grounds that Kitcher's analysis of a priori knowledge is flawed, and the failure of mathematical knowledge to satisfy (2) is thus irrelevant to its a priori status. This would be a poor move for at least three reasons. Firstly, Kitcher claims that his is the only thoroughly developed notion of a priori knowledge in the literature—reject it, and the defender of the a priori has no clear concept to defend. Regardless of the truth of Kitcher's claim here, it is fair to say that it is incumbent upon one who rejects Kitcher's analysis to provide an alternative. In the course of my treatment of Kitcher's arguments, an account of a priori knowledge which is based on Kitcher's account, but which does not have the consequences he claims for his own account will emerge.

Secondly, at least some of what Kitcher has to say is independent of the nuts and bolts of his analysis, and, more importantly, where this is not so, close attention to his arguments can serve to refine our own conception of a priori knowledge.

Finally, one might be tempted to reject (2) because clause (c) requires that if one is warranted a priori in believing that p, then p. One might, that is, reject (c) (and hence (2)) because it requires of a priori warrants that they be infallible. It is true that we should not expect of a priori warrants that they be infallible, and if Kitcher's account relied crucially on this, it would be fatally flawed. However, as Hale[89] observes, Kitcher's arguments rest rather on clause (b). And clause (b) is not so easily rejected. Kitcher (1983, 98-9) had argued that to reject (b) is to abandon the idea that a priori knowledge is independent of experience:

> The apriorist would be saying that one can know a priori that p in a particular way, even though, given appropriate experiences, one would not be able to know that p in the same way. But if alternative experiences could undermine one's knowledge then there are features of one's current experience which are relevant to the knowledge, namely those features whose absence would change the current experience into subversive experience ... To reject condition (b) ... would be to strip apriorism of its distinctive claim.

[89] Bob Hale, Abstract Objects (Oxford: Blackwell, 1988), 129.

However, Kitcher (2000, 73) has moved away from this position, citing the following example of Charles Parsons: The fact that I can imagine having grounds for doubting my current perceptions does not imply that the absence of those experiences are currently sustaining/generating those beliefs. I thus do not have to establish that I am not the victim of a Cartesian demon every time I perceive an object in order to be warranted in my visual beliefs.[90] Thus, Kitcher concedes, the fact that there are conceivable features, which, though absent, would undermine my warrant for a belief that p, were they present, does not entail that the absence of such features is currently warranting my belief that p.

However, even without Kitcher's back-up argument, there is reason to accept (b). Clause (b) demands of an a priori warranting process that it would warrant p when it produces or sustains belief in p. Now, to say that a process would warrant a belief is not to say that the process invariably warrants belief in p any time it occurs, but to say that across a range of relevant possible worlds, when that process yields belief in p, p is warranted. Clause (b) thus imposes only a reliability condition on a priori warranting processes. As a result, (b) seems mandatory for the reliabilist, and, as we shall see, comes to prove quite amenable to the apriorist.

Since the conception of a priori warrant to be defended here is fallibilist, I do not accept Kitcher's condition (c). (c), as Bob Hale observes, is a truth condition, rather than a condition on warrant. I thus propose to examine the question of whether mathematical proof correctly counts as an a priori warranting process by showing that it

[90] The Parsons example is from his "Review of The Nature of Mathematical Knowledge," Philosophical Review 95 (1986), 129-37.

satisfies Kitcher's analysis, when clause (c) is dropped, and clause (b) is understood in what I take to be the most plausible way. [91] Rejecting (c) leaves the way open for a warranting process to yield warranted beliefs which are in fact false.[92] Apriorists, then, need not—and indeed, had better not—commit to (c). However, as I have already argued, they cannot avoid Kitcher's main arguments by simply rejecting his entire analysis as fatally flawed due to the presence of (c). One who is both a mathematical apriorist and a reliabilist would do well to examine Kitcher's arguments in detail.

### 3.1.1

### Apriorism and Rational Uncertainty (i): The Decay Picture

Kitcher's (1983, 40) first point against mathematical proof as an a priori warranting process rests on a worry about long proofs. Here is Kitcher:

> There are true standard mathematical statements S such that the shortest proof of S would require even the most talented human mathematicians to spend months in concentrated effort to follow it. Can we really suppose that S is knowable a priori? After all, anyone who had followed a proof of S would reasonably believe that he might have made a mistake ... So we might conclude that our knowledge of S is inevitably uncertain, and therefore not a priori.

It is here, especially given the closing sentence in the quoted passage, very tempting to simply deny that a priori warrant must imply certainty, and thus reject the argument.

---

[91] There is, however, an obstacle in the way of this proposed method of proceeding. Kitcher (2000, 72) argues that clause (c) is a consequence of the other clauses, and thus that such cherry-picking as I am proposing cannot be defended. This argument will be taken up in section 3.1.4.

[92] As we shall see, however, it does not leave open the way for an a priori warranting process to become non-warranting. If this were possible for an a priori process, it would violate (b), and would thus, according to the conditions I am accepting from Kitcher, cease to be an a priori warranting process.

But things are not that simple. The worry about long proofs is exacerbated only a few paragraphs later, when Kitcher considers what he calls the "Decay Picture" of a priori warrant. Perhaps when we follow a long proof, what goes on is that the conclusion from each inference possesses a degree of warrant slightly less than that of each of the premisses. Given long enough proofs, then, the Decay Picture suggests that at least some mathematical statements will not posses enough warrant to be knowable a priori.

The question the apriorist ought to ask at this point is, of course, why should we think that the process of following a proof—even a very long one—is correctly described by the Decay Picture. Kitcher (1983, 41) suggests that support for this model is to be found in Hume's observation that no mathematician is so sure of his ability as to be completely confident in a truth he has just discovered. Rather, according to Hume, his confidence is raised by (a) retracing his steps, and (b) having his proof approved by fellow mathematicians.

The support which Kitcher claims from this example would, presumably, go something like the following. If the proof were a simple one, there would be no need to recheck it, nor to offer it to others to do so. In following the proof, one would have all the warrant one needed. However, once the proof exceeds a certain length or complexity, one is naturally less confident in the conclusion. In order to become sure of the proof, one must recheck it, and perhaps offer it to colleagues for their perusal; this indicates that one has less warrant for believing the conclusion. The Decay Picture predicts this, and so gains support from the existence of such cases.

One immediate response available to the apriorist here is to deny that Hume's description of the psychological state of a mathematician who has discovered a new proof has any epistemological import. Kitcher (1983, 42) anticipates such a move, claiming that there are only two reasons for denying the epistemological import of the psychological state. The first is to claim that the uncertainty is an accidental feature of the situation, stemming from the incomplete nature of proofs. The second is to deny that a priori warrant implies certainty. I shall, for the sake of argument, simply grant Kitcher his response to the first point here.[93] My argument for denying that the psychological state of the mathematician has any epistemological import, then, will rest on an examination of the relationship between a priori warrant and certainty. However, before taking up this issue, I wish to point out another response which can be made in response to Hume's observation. The digression is worthwhile, both because it uncovers a point Kitcher appears to have overlooked, and because it ultimately leads back to the issue of the relationship between certainty and a priori warrant.

It is a commonplace in the epistemological literature[94] to note that belief, like warrant, comes in degrees. If anything like this picture is true, then it seems that, in order for a belief that p to be knowledge that p, one would have to have something close

[93] The basic thrust of which is to note that if we are impressed with the decrease in the confidence of mathematicians when it comes to long proofs, requiring that we make proofs "more complete" by insisting on stricter standards of formal rigor will not ease our worries. Such insistence would merely result in longer proofs, more room for error, and thus, even less confidence.

[94] Examples abound, but to note just two, see Bas Van Fraassen "Belief and the Will," Journal of Philosophy (1984): 235-56, and Richard Foley "How Should Future Opinion Affect Current Opinion," Philosophy and Phenomenological Research (1993): 747-66.

to full belief that p.[95] The apriorist, taking this observation on board, can deny that the case of the uncertain mathematician provides any support for the claim that the process of following a mathematical proof is suitably modeled on the Decay Picture simply by considering the uncertainty of the mathematician. If the mathematician is very uncertain, then he does not know the conclusion in question. But this is not because following a long proof fails to confer an a priori warrant. It is rather, because the mathematician does not really believe the conclusion. He is in a state of partial belief, at best, or of not believing at all, at worst. On the assumption that the proof is reliable and leads to a true p, the mathematician will come to know that p once he (fully) believes it. The reason that the mathematician in question doesn't know the proof's conclusion a priori has nothing to do with mathematical truth not being a priori, but rather, has to do with his failure to satisfy the belief condition on knowledge.

Suppose, however, that it is insisted that the mathematician both (a) believes the conclusion in question (to a level sufficient to meet the condition on knowledge), and (b) is uncertain of it. The question is, why does this rule out his knowing the conclusion a priori? After all, he believes it, it is true, and the proof he followed was reliable; why not just say he does know it, and that he knows it a priori? Kitcher (1983, 43), explains matters thus:

> Reasonable uncertainty is typically compatible with knowledge because of the kindly nature of background experience. Transform the quality of our lives, and

[95] Just how close one's belief would have to be to full belief (or, how close to 1 one's "subjective probability" would have to be), I here leave open. The important point is that if one's belief is substantially less than full (if one's subjective probability is substantially <1), then one would not count as knowing that p, simply because one would not believe it sufficiently.

> that knowledge could no longer coexist with the uncertainty ... Rational uncertainty does not preclude knowledge but it does rule out a priori knowledge.

The crucial question, then, is why rational uncertainty is supposed to rule out a priori knowledge. According to Kitcher, this is because condition (2b) is not satisfied by processes which produce mathematical beliefs. Take a mathematician who has just followed a proof and come (correctly and on the basis of this proof) to believe its conclusion. Call this scenario A. Now, take that same mathematician, forming that same belief by a process of the same type, and change his environment. Specifically, place him in an environment where most other reputable mathematicians tell him that he has made an error in his reasoning. Call this scenario B. The mathematician in scenario B, according to Kitcher, would be unreasonably arrogant to continue in his belief. That is, the process which had produced warrant for his belief can occur without conferring warrant. Condition (2b) is not satisfied, so the mathematical truth in question is not known a priori.[96]

Condition (2b), recall, requires that if a process of the same type (as one which allegedly confers a priori warrant on the belief that p) were to produce in X a belief that p, then it would warrant X in believing p. If Kitcher's case is to be a counterexample to

[96] Although I remarked earlier that I would not dismiss Kitcher's arguments on the grounds that they employ a fundamentally mistaken conception of the a priori, it is important to note that his conception of a priori warrant as implying indefeasibility is extremely strict. Casullo notes that, since empirical warrant is not supposed to be indefeasible, to impose a stronger condition (i.e. indefeasibility) on a priori warrant is simply ad hoc (see his "A Priori Knowledge Appraised," in A Priori Knowledge, ed. by Casullo, Aldershot: Dartmouth Publishing Company, 1999). And certainly there are accounts of a priori warrant which allow defeasibility (Katz, 1998), and even some which allow empirical defeasibility (See Tyler Burge, "Computer Proof, A Priori Knowledge, and Other Minds," in Philosophical Perspectives, 12, Language, Mind and Ontology (Boston, Mass.: Blackwell, 1998)). So it certainly seems that Kitcher's account imposes very strong conditions on a priori warrant. My argument, however, will be that even given Kitcher's conditions (2a) and (2b), there is no reason to think that a priori warranting mathematical processes can occur without conferring warrant.

mathematical apriorism, then, it needs to show both (i) that the same type of process is at work in both cases, and (ii) that it confers warrant in one case, but not in the other. In fact, Kitcher's example shows neither of these.

Consider (i). The mathematician in scenario A comes to believe that p by following a proof. As does the mathematician in scenario B. So far so good for Kitcher's argument. Things begin to unravel, however, when we notice that, at this point in the proceedings, the mathematician is equally warranted in both scenarios. That is, the difference in degree of warrant does not arise in connection with the belief-producing processes. It arises rather, in connection with the belief-sustaining processes. Moreover, this difference in degree of warrant is explained by the fact that the processes which sustain the belief in each case are completely different.

The mathematician in scenario A comes to believe that p by following a proof, and continues (with accessible warrant) believing that p through remembering that he has followed a proof (perhaps also by running over it again) and through not encountering any reason for doubting that p. The mathematician in scenario B, on the other hand, continues in his p-belief in spite of having very good reasons to doubt p. Many (perhaps all) of his professional colleagues tell him that the proof is faulty, that he has made a mistake, that he has overlooked something, and so on. His p-belief is sustained in the face of very weighty considerations for doubting p. But if we remember that processes are not just belief-producing, but also belief-sustaining, then it seems clear that the two processes are not at all similar. The process in scenario A is that of believing p on the basis of following a proof sanctioned by all reputable mathematicians (and where the proof is in fact reliable). The process in scenario B is that of believing p on the basis of following a proof (which, although in fact sound, is) rejected by all—or at least many—reputable mathematicians.

Now it might here be objected that if we see processes as being individuated purely internally, as Kitcher does, then we have to say the processes which sustain the beliefs of the mathematician in scenarios A and B are the same.[97] But this is not so. For, in scenario B, the mathematician will have heard reasons why he has to give up his p-belief. If he considers these reasons, then his continuing to hold that p will involve a process which either discounts these reasons, or holds on to beliefs in the face of compelling reasons to the contrary. Clearly neither of these processes are at work in scenario A. If he refuses to consider these reasons, then the process which sustains his p-belief in scenario B will involve deliberately ignoring evidence in his possession. Again, no such process is involved in scenario A.

Kitcher has objected to this line of argument that such fine-grained individuation of processes will make condition (a) harder to satisfy, and so the process will fail to be a priori in a different way.[98] There is, as the generality problem illustrates, a taxonomy problem for the reliabilist. In Chapter One, I suggested that the most obvious way of individuating logico-mathematical processes was to use the already-present distinctions between different rules of inference, or proof-procedures to demarcate different processes. Thus, using the same rule of inference on different occasions would count as

[97] Thanks to Michael Levin for raising this objection.

[98] The objection was made by Kitcher in response to an earlier (and much shorter) version of this chapter, which I presented to an audience at Columbia University.

using the same process on different occasions. This method, however, only works for simple cases where a belief is generated or sustained by one process. Where more than one process generates or sustains a given belief, then, in order for us to correctly say that the same process generates (or sustains) this belief in a different situation, all the relevant generating (and sustaining) processes must be present. This does make the question of which is the correct description of a given (set of) process(es) harder to answer, but for present purposes, there seems an obvious difference in the two cases. I suggest that (part of) the sustaining process for this belief is clearly that of ignoring evidence which one has reason to think relevant to evaluating p. The presence of this process (IG) is enough to generate a difference between the two cases, so Kitcher's point does not go through. Nor does it seem at all obvious to me that slicing processes on such a commonsensical level represents slicing processes so fine grained that we could not satisfy (a). In the present instance, (a) would be satisfied if the additional process IG were present in both cases. In any event, it certainly would need to be shown, and thus far it has not been, that my move renders (a) unsatisfiable. Thus, Kitcher's objection to my point that condition (i) has not been satisfied does not show what it needs to show, and the point stands. The case is not a counterexample to mathematical apriorism, because (i) is not satisfied by it.

Consider now, (ii) above. Suppose Kitcher is correct in claiming that a process of the same type produces our mathematician's p-beliefs in both scenarios. Is it correct to say that his p-belief in scenario A is warranted, but his p-belief in scenario B is not? As Hale (1988, 135) remarks, if Kitcher's argument is to show what it seeks to show, we must assume that the proof which we have followed is, in fact, sound, and that the experiences which undermine this proof are misleading (otherwise, the belief in the proof's conclusion was never knowledge, and, a fortiori, never a priori knowledge). On this assumption then, together with the assumption (which nothing so far has called into question) that following a proof is generally a reliable belief-forming process, the apriorist has a straightforward response to Kitcher. As the particular proof has been correctly followed, and has led to the formation of a true belief, and as such processes are completely reliable, our belief in the conclusion of the proof is an instance of knowledge a priori. If our mathematician has correctly followed and understood the proof, he knows that any considerations to the contrary are mistaken. Unkind experience might offer what appeared to be potential defeaters for this belief, but since the method is (completely) reliable there could be no actual (or indeed possible) defeaters. The fact that one does not consider an apparent defeater to a proposition which cannot have any defeaters is powerless to affect one's knowledge. To put things in more specifically reliabilist terminology, the fact that one did not consider ways in which the process leading to one's belief might have been shown unreliable, when in fact such a process could not have been shown to be unreliable cannot affect the reliability of that process. True mathematical beliefs, formed by the process of following a proof, are instances of knowledge a priori.

Of course, Kitcher would not accept the conception of mathematical proof as being invulnerable to being shown unreliable. But this way of understanding proof is a central part of the Platonistic brand of apriorism here being defended. According to this position, a proof is an abstract, unchanging structure, the premisses of which necessarily entail the conclusion. Correctly following a proof, on this picture, means instantiating a mental token which accurately represents that abstract structure. Thus the process of correctly following a sound proof is one which cannot have any defeaters, and thus cannot be shown unreliable. Kitcher's cases might work against non-Platonist apriorist positions, but if he is to argue successfully against the apriority of mathematical proof, then he needs to do so in a way which does not simply assume the falsity of central aspects of one of the main apriorist positions.

One might question the claim that not considering apparent defeaters to processes which couldn't have defeaters is powerless to affect one's knowledge. If most of the mathematical community were against you, continuing to ignore their reasons would not usually be a reliable method. This would not affect the main point, however, since here the reasons, whatever they are, are, ex hypothesi, misleading. In such cases, refusing to consider the mistaken reasoning in opposition to the sound proof is powerless to affect one's knowledge. And, it bears repeating, that by reliabilist lights, our mathematician's belief, as it is both true and produced by an ultra-reliable method, is an instance not only of warranted belief, but of knowledge. This point is compatible with the fact that, in general, ignoring the opinions of most other mathematicians would not be a reliable method, since there would be a good chance that they would be right. Here, however, the unreliability of this method is offset by the ultra-reliability of the process of correctly following a sound proof.

A further objection arises from the consideration that mathematicians, prior to the discovery of the set-theory paradoxes would probably have believed that no defeaters were possible for the belief that every class determines a set. As they were mistaken, what justifies the response given to Kitcher's example? The answer is that we are assuming that the proof followed is in fact sound, and this differentiates the case from that of pre-paradox set-theorists. They, of course, were not aware of this difference, but this does not affect the point that it is a relevant difference in the cases.

One final objection to the line of defense pursued here, and which picks things up where the previous objection left off, is that the mathematician in question, although following a completely reliable process, does not know that he is, unless he can rule out the possibility of further a priori considerations which show that his proof procedures are unreliable. There are two responses to this point. Firstly, once one abandons the equation of a priori warrant with indefeasibility, there is no motivation for requiring that one rule out all imaginable defeaters before one can be said to know. This point is generally acknowledged with respect to empirical beliefs (such as my belief that a particular bird is a bullfinch); it should be insisted upon with respect to a priori beliefs. Secondly, given the reliabilist outlook of this dissertation, one can respond that if one is in fact correctly following a completely reliable proof, then one has a priori knowledge. The fact that one is not in a position to rule out future a priori discoveries which would undermine one's proof procedures would be relevant to the question of whether one knows that one knows the conclusion of the proof, but not to the question of whether one simply knows it.

The case of unkind experience, then, shows neither (i) nor (ii). Kitcher has therefore failed to present a counterexample to mathematical apriorism. The case was supposed to show that rational uncertainty was incompatible with a priori warrant. This, in turn, was supposed to show that the uncertainty of mathematicians concerning long proofs had epistemological import. And this last was supposed to support the claim that the following of long proofs was accurately depicted by the Decay Picture. As Kitcher's case does not show what it needed to, the attempt to support the Decay Picture fails.

### 3.1.2

### Apriorism and Rational Uncertainty (ii): Methodological Disputes

One might object to the foregoing, as Kitcher has,[99] that the above case is too easy for the apriorist, focusing as it does on a simple case of a mathematician in possession of a sound proof, and where methodological principles are not in dispute. Such a picture makes it appear simply obvious that the apriorist is correct, and that the discordant mathematicians are being little more than obtuse. More difficult would be cases where the dispute is over methodological concerns, so that the outcome of such disputes would determine which kinds of procedures are to be counted as reliable. Consider the following kinds of cases:

Type 1: Cases where proposed methodologies are rejected because they eventually prove to lead to contradiction.

Type 2: Cases where new procedures are not adopted simply because the mathematician who has proposed them lacks the required influence to "sell" the new procedures to his colleagues.

Type 3: Cases where proposed methodologies are rejected because, although otherwise unobjectionable, they are not fruitful.

---

[99] Kitcher and Federico Manrulanda Rey both raised this objection when I presented part of this chapter as a paper at the 2001 Columbia Graduate Student Conference on Rationality. The type-four case is due to the latter. The remaining case were all explicitly raised by Kitcher at that time. My thanks to both.

Type 4: Cases where the debate is over whether to extend a concept beyond its previous applications.

Type 5: Cases where a putative proof is sound with respect to one mathematician's proposed proof-procedures, but not with respect to the proof-procedures of his colleagues.

What each of the above cases has in common is that the standards for reliability seem to vary across the mathematical community. In each of Type 1-5, a particular way of doing mathematics will be accepted by one (set of) mathematician(s), and rejected by others. Sometimes the disputant mathematicians will belong to the same generation, and sometimes, as in the case of procedures which are seen to lead to contradiction only after a long time, the methodological dispute will be between ages. In either event, there will be a situation where, according to those who accept some controversial proposed process (e.g. a particular new proof procedure) this process will count as a priori warranting, but amongst mathematicians who do not accept the proposed process, it will not be so counted. So, the argument goes, condition (b) is not satisfied, and non-basic mathematical knowledge does not count as a priori. From this perspective, the fact that standards for reliability within mathematics seem to vary in this way shows what was wrong with the above response to the argument from unkind experience. That response depended on the claim that mathematical proofs were completely reliable. Since these harder cases are ones where there is dispute over which kinds of proof procedures are sound, they appear to show that this is not so. The response seems, therefore, to beg the question against Kitcher.

Type-1 cases are different from the others, in that they are not ones where there is a mere difference of opinion as to whether a given process is warranting. Instead they are cases where a process has been demonstrated to lead to contradiction, a result which everyone agrees is unacceptable. I shall deal with Type-1 cases separately, and then with the others.

Let us take a real-life example, that of Frege's principle that every property determines a set being shown to lead to contradiction. It will prove useful to recall the distinction between a warranted belief, and a warranting process. A warranted belief, according to the reliabilist, is simply one that was produced by a reliable process. There is room, therefore, for a belief which is warranted, but false; all that is required is that the false belief in question be produced by a (usually) reliable process. A warranting process, on the other hand, is a process which confers warrant on beliefs—beliefs either produced or sustained by that process. In the present case we are concerned not with whether the belief that every property determines a set is warranted. This would be relevant to the question of the reliability, or fallibility, of the process which produced that belief. In Chapter Five, I shall argue that this belief, though false, was warranted. No, we are here seeking a case which shows that a prima facie a priori mathematical process violates (b). We are therefore concerned with whether the process of using this principle to further mathematical research is an a priori warranting process. The relevant question, then, is whether the case is one in which a genuinely warranting, a priori process becomes non-warranting. The answer, it should be obvious, is that it is not. The process has been demonstrated to lead to contradiction. Prior to the discovery of the paradoxes, the process still implied a contradiction, even though that contradiction was not yet discovered. Using the process was thus never reliable, even though it was at one time believed to be. Since it never was warranting, the fact that it is not now warranting does not give us a case where a warranting process becomes non-warranting.

Still, assuming that the belief that every property determines a set was itself produced by a reliable process (for example, by mathematical intuition), an a priori warranted belief has had its warrant defeated. Does this represent a problem for the apriorist? (b) requires of a putative a priori warranting process that if it were to produce in a subject the belief that p, then it would warrant the subject in believing that p. This idea gets its credibility from the idea that if a belief is warranted a priori, it should not be hostage to future experience. If it is in fact a priori, then we ought to be warranted in believing it no matter what experiences we have. Indeed Kitcher's introduction of his account of a priori warrant makes it clear that he is leaning heavily on this intuition. In a discussion of (c), Kitcher (1983, 24) states that his goal is "to construe a priori knowledge as knowledge which is independent of experience" (my emphasis), and tells us that "to generate knowledge independently of experience, a priori warrants must produce warranted true belief in counterfactual situations where experiences are different" (final emphasis mine).

But there is an ambiguity in the notion of experience. On the one hand, experience means something like 'information deriving from the senses, which is, by nature, contingent and a posteriori, the truth of which cannot be ascertained a priori'. I shall refer to experience of this kind as "narrow experience." I take it that it is experience in this sense that is relevant to the debate between rationalists and empiricists. However, there is another sense of 'experience', which I will refer to as "broad experience," which Kitcher makes much of, according to which any subjective mental occurrence counts as experience. On this conception, any conscious intellectual event, even understanding a proof, or reading about a methodological debate in mathematics counts as experience. That Kitcher's understanding of 'experience' incorporates what I am calling broad experience is clear in both of the works under discussion. In The Nature of Mathematical Knowledge, addressing the question of whether there can be experiences which suggest the falsity of mathematical statements, Kitcher (1983, 64) claims that there are three kinds of experience which can: social challenges, direct challenges, and theoretical challenges. Certainly, without going into the particulars of what the latter two involve exactly, it is clear that these two, if not all three, go beyond the scope of narrow experience. Likewise, in "A Priori Knowledge Revisited," Kitcher (87) imagines "experiences" which might suggest that he is mistaken to believe some laws of classical logic. The kind of experience envisaged is that of being shown that the laws in question don't correctly reconstruct the inferences that are made in those parts of knowledge that we would be most reluctant to revise.

But the idea that what is a priori warranted should not be vulnerable to future experience is only credible where future experience is experience in the narrow sense. According to the fallibilist picture, an a priori warranted belief is, by definition, vulnerable to being undermined by future a priori considerations, and thus, by future broad experience. How this happens is nicely illustrated by Russell's discovery of the paradoxes, and the revisions this caused in Frege's set theory. The whole process is fallible a priori thinking correcting itself. At no point does anything like narrow experiential warrant come into the picture. The paradoxes were, after all, discovered a priori, without reliance upon the senses. The problems they presented in the foundations of mathematics were understood, and suggestions for dealing with them offered. The discovery of the set-theoretical paradoxes thus has no tendency to undermine the apriorist's position.

Of course, one could, if one wished, use the broad sense of experience, and describe this event as Frege's "experience" of Russell's paradox occasioning revision of an a priori warranted belief, but describing things in this way does not threaten mathematical apriorism. A threat would arise only if experience in the narrow sense entered the picture, and rendered an a priori warranted belief unwarranted. This is because if broad experience is what is meant by experience, then mathematical empiricism is vacuously true. If "experiences" of understanding refutations of supposed theorems count as experience, then the concept of beliefs whose warrant is not dependent on experience is not one worth defending.[100]

Moreover, if we are to understand by 'experience' narrow-experience, we see that the kinds of cases which are relevant to the question of whether (b) is satisfied are cases where the only elements that change between two sets of circumstances are elements of narrow experience. Cases where more than narrow experience changes between two sets of circumstances (for example, where one has the broad experience of reading a refutation of a particular mathematical belief) are not relevant to the assessment of whether an allegedly a priori process has conferred experience-independent warrant on a given belief. Only those worlds in which everything other

[100] Tyler Burge suggests that Leibniz at times seems to understand 'experience' in this broad way. See his "Frege on Apriority," in Boghossian and Peacocke, (2000, 28).

than narrow experience is held constant are germane to the question of the satisfaction of (b) by a given case.

However, even given the distinction between narrow and broad experience, doesn't the Frege case show that some prima facie a priori warranting processes can, in some circumstances, occur without it warranting p? And doesn't this show (b) to be violated by such processes? No. Cases where an a priori warranted belief has its warrant undermined by broad experience are not cases which fail to satisfy (b). This is because a belief can be unwarranted due to the sum total of relevant processes warranting not-p. This is not to say that one of the set of processes that warranted p has failed to confer warrant on p. The process in question still confers (and would confer) warrant on p, but that warrant has been undermined by the warrant for not-p conferred by the sum of the other relevant processes. Thus, (b) is satisfied, although belief in p is not, on balance, warranted. The belief that every property determines a set, then, was false-but-warranted. In fact the same belief, produced by the same (ex hypothesi) reliable process would still have some warrant conferred on it by the producing process. This can be seen by imagining an auto-didact mathematician, unaware of the paradoxes, in whom the belief, formed by a reliable process, that every property determines a set, has just been produced. His belief would be warranted. Those of us who are aware of the paradoxes would not, on balance, be warranted in believing that every property determines a set. But this is not because the producing process no longer confers warrant. The process is still reliable, even if it has produced a false belief, and it is thereby warranting. No, our belief is, on balance, unwarranted because there are other, more reliable processes, which show that belief to be false

Type-1 cases, then, do not present a problem for the fallibilist apriorist. What about the remaining case-types? Are they cases in which a prima facie a priori mathematical process is shown to lose its warranting power, and is thus shown not to be an a priori warranting process? The first thing to be said in reply to this is that, if this is really what is going on in such cases, it shows at most that controversial elements of mathematics are not a priori. In spite of the presence of such methodological disagreements in the history of mathematics, there is still an overwhelming amount of mathematics where methodological principles are not in serious dispute. Even if Kitcher were correct, and those parts of mathematics which involve methodological disagreement were thereby debarred from the a priori, the apriorist would still have arithmetic, simple geometries, trigonometry, algebra and calculus to lean on, to name only the most obvious examples. For all Kitcher has said, these would appear to comfortably satisfy both (a) and (b), and thus count as a priori. The apriorist would thus, still emerge victorious, albeit with a somewhat trimmed-down range of a priori mathematics.[101]

However, I do not think that Kitcher's point is correct, even when limited to controversial areas of mathematics. I shall argue that none of case types 2-5 are genuinely types of cases where (b) is violated. In arguing this, I shall make use both of the Platonist conception of mathematical objectivity, and the reliabilist conception of warrant.

What each of the remaining case-types have in common is that there is disagreement about whether a given process is reliable. Now, disagreement on its own

[101] Thanks to a member of the audience at Columbia University for suggesting this response.

does not establish anything about whether a process confers warrant in a given situation, but not in another. In order to show that (b) is violated, what is needed is a case where a particular process in fact confers warrant in one set of circumstances, and in fact fails to confer warrant in another set of circumstances. These harder cases are not of this sort, for they are only cases where mathematicians differ in opinion about whether a process is reliable. In order to convert this kind of disagreement over methodological issues into a case which shows a prima facie a priori mathematical process to fail (b), it must further be assumed that <u>what mathematicians believe about a process</u> is part of what confers warranting powers on a process. As this assumption runs contrary to the Platonist's understanding of mathematical objectivity, it is not one he is likely to grant.

If, contrary to what I am urging, the ability of a mathematical process to confer warrant on a belief were dependent on which other processes and principles were accepted by the mathematical community, we can see that (b) would be unsatisfied by the above cases. If the reliability of some process, r, were dependent on the fact that it followed from some set of axioms, A, which was accepted by the mathematical community, and if A were subsequently to be rejected, then r would become unwarranted. Under such circumstances, we would have a warranting process becoming non-waranting, and (b) would be violated. However, for one who holds that the warranting powers of processes are independent of our opinion, this kind of case does not show anything about (b).

To use a specific example, the fact that two sets of mathematicians are at odds about whether (e.g.) the Law of Excluded Middle is a reliable process does not show that the LEM is unreliable were the intuitionists to convince everyone that they were correct, and reliable in the different set of circumstances where mainstream mathematicians do so. It shows only that the opinion of mathematicians as to its reliability differs in these circumstances. If the LEM is actually warrant conferring, then it would be so in both sets of circumstances. If, in addition to being reliable, it happens to be rejected, this does not show that there are circumstances where it does not confer warrant (and thus that (b) fails), but only that there are circumstances where it is <u>thought</u> not to confer warrant. This does not show (b) to fail, and thus, even on Kitcher's own account of the a priori, these hard cases do not show what they need to show to worry the apriorist.

Similarly, the fact that an a priori warranting process does not become generally accepted due to the lack of influence of its discoverer shows no more than that some actually warranting processes slip through the net. It does not show that a prima facie a priori warranting process can, under some circumstances, fail to confer warrant. And, as there is no strong commitment on the part of the apriorist to mathematicians having the unerring ability to pick out every single warranting process, the disagreement here is likewise harmless.

We have an answer, then, to the claims, made on behalf of Kitcher above, that standards of reliability for mathematical procedures shift across circumstances, and thus that to assume that mathematical procedures are completely reliable begs the question against Kitcher. What shifts is not the standards for the reliability of a priori processes, but our standards <u>for assessing the reliability</u> of such processes. A reliable a priori mathematical process is a reliable process, whether we recognize it as such or not. The soundness of sound proof-procedures and the validity of valid methodological principles thus do not depend on our opinion. This is part of the Platonist's position. There is mathematical truth, independent of us, and we discover it. Disagreements according to this picture, are disagreements about independent objective reality. They are powerless to change that reality, and thus cannot change what was once a warranting process into one that no longer confers warrant.[102] It is true that we will sometimes answer the question of whether a particular procedure is reliable incorrectly. Likewise, it is true that sometimes (seemingly more often, if the success of mathematics is any indication) we will get it right. We will miss out on some methods, and hit on others, and this itself will help determine which further methods we deem reliable. But none of this has any tendency to show that standards for the reliability of a priori mathematical processes shift across circumstances.

One might object that the foregoing response confuses (the Platonist conception of) independently existing mathematical process with the mental process of following such a process. Since we are interested in beliefs and believers, it is the latter that are relevant, and since these involve change from one contingent physical state to another, they, unlike their eternal unchanging correspondents, can warrant belief in some circumstances, and yet fail to do so in others. Thus, (b) is violated after all. The objection fails, however, because if the mental process in question successfully instantiates an actually warranting process, then it will warrant belief. A putative metal instantiation of an a priori warranting process could only fail to warrant belief if in fact it failed to instantiate such a process.

[102] This point is not in conflict with the central fallibilist apriorist claim that one can have false-but-warranted beliefs. On the line being sketched above in the text, false-but-warranted beliefs are precisely those (false) beliefs which are produced by in-fact warrant-conferring processes.

A second objection to the proposed line arises from consideration of how a posteriori processes which are actually warranting, might become non-warranting. Take, as an example of an a posteriori process, the visual process. On earth, this process is fairly reliable. However, we can imagine circumstances in which a human were to be placed on a planet where vision were to regularly yield unreliable information, due perhaps to the obtaining of optical laws different from those obtaining on earth. Under such circumstances, we would have an in-fact reliable process becoming unreliable (and thus, a warranting process becoming non-warranting), and this would be due to a change in the context in which that process is embedded. Now Kitcher is attempting to create the same possibility with respect to mathematical processes. That is, he is attempting to show that a change in the embedding context in which an in-fact warranting mathematical process is used, can render that process non-warranting. And even though the relevant embedding context here is intellectual, rather than sensory, the basic move is the same. The reliabilist-Platonist response to this move is to say that this kind of case cannot arise for genuinely a priori mathematical processes.[103] And the question is 'Why not?'

One might be tempted to say that the reason that prima facie a priori mathematical processes cannot become non-warranting is because the propositions in which such processes yield belief are necessary truths. This line has it that since mathematical truths are necessarily true, and since a process which produces belief in necessary truths will have its reliability guaranteed by the nature of the beliefs it

---

[103] Kitcher of course agrees with this point, but only because he believes that there are no genuinely a priori mathematical processes.

produces, there is no room for the process to become non-warranting. But, as we saw in a different context in Chapter One, this response does not quite work. Since there might be processes which produce belief only in propositions the truth or falsity of which is necessary, but which regularly produced belief in necessary falsehoods, the mere fact that a process produces belief only in propositions which have their truth value necessarily would not be enough to insulate such processes from the possibility of change in their embedding context rendering them non-warranting. Even though such a process traffics in propositions that have their truth values necessarily, it does not reliably track the truth of such propositions. There can be processes which produce mathematical beliefs, which can change from being warranting to being non-warranting. My honest mathematics teacher can become a liar, and thus the process of obtaining my beliefs from him will change from being a warranting process, to being a non-warranting process. Which is to say, it would be shown not to be an a priori process.

As one who is both reliabilist and Platonist sees things, a priori mathematical processes reliably (though fallibly in the cases of some such processes) track the truth of necessary mathematical propositions. Because the truth values of mathematical propositions do not change, if a process reliably tracks these truths, its warranting status likewise cannot change. In order to reliably track the truth of necessary truths, a given process must be able to track these truths regardless of changes in the embedding context. The Platonist's claim is that prima facie a priori mathematical processes, such as using intuition, and following proofs, are genuinely a priori since they do reliably track the truth of necessary truths.[104] If the ability of a process to track such truths can be made unreliable by changing something in the embedding context, this cannot be because the truths have changed, since they cannot. It simply means that the reliability of the process was dependent on the embedding context in which it was used, and thus that the process was not a priori warranting.[105] Since the Platonist holds that a priori processes are ones whose reliability is context-independent, he will hold that if Kitcher can give us a case where the warranting status of a mathematical process changes, that simply means that the process was never an a priori warranting process to begin with. Of course, Kitcher's claim is precisely that he is showing the warranting powers of prima facie a priori mathematical processes to be vulnerable to changes in the embedding context of use. But I have already argued that disagreement about the warranting status of a process does not imply change in the warranting status of a process. Nor has it been shown that the reliability of prima facie a priori mathematical processes is vulnerable to change in the embedding context. By Platonist lights, this cannot be shown, because tokening such processes instantiates mind-independent abstract sequences the premisses of which necessarily confer warrant on the conclusion. Change in the embedding context in which the process is used would be change in the subject's intellectual or experiential setting. By the same argument from Platonist

[104] This is not to close the door on the possibility of there being processes which reliably track contingent a priori truths. Reliably tracking the truth of necessary mathematical truths is only one way in which one could come to have a priori knowledge. It is not a necessary condition on all a priori knowledge, just on knowledge of a priori necessary truths.

[105] That a priori mathematical processes reliably track necessary mathematical truths, whereas empirically warranting processes do not, explains why a priori warrant can undermine, and is not open to being undermined by, empirical warrant. A priori processes are ultra-reliable, in the sense that they reliably track necessary truths. Empirically warranting processes, since their warrant can be undermined by change in the embedding context, do not reliably track such truths.

ontology given above, if a proof fails to warrant a mathematical belief, that can only be because that process does not in fact instantiate the relevant proof.

The Platonist's corresponding claim with respect to dispute over methodology, is that the correctness of any given answer is settled by its corresponding to, or failing to correspond to independent mathematical reality. Take for example Cantor's demonstration that there are more real numbers than there are natural numbers. This is a surprising result; the conception of a hierarchy of infinities to which it leads is still not accepted by all. The decision to accept the result involves an extension of mathematical concepts, so as to allow for the difference in numerosity between real numbers and natural numbers. But the concepts might have been extended in a different direction, so as to block this result. The Platonist holds that, for any such extension of mathematical concepts, the decision to extend in a particular way either corresponds to mathematical reality, or it does not. Our opinion here does not alter the fact of the matter.

We can have confidence in a given instance of an extension of concepts on the basis of how fruitful that extension is, its robustness in the face of attempts to show it leads to contradiction, the fit between it (and its consequences) with the rest of accepted mathematics, and how well it allows us to solve other problems in mathematics. All of which give us reason to believe (fallibly) both that we have extended our concepts in the right way, and that the procedures which have led to their acceptance reliably track necessary truths in a context-independent manner (and are thus warranting). We could be mistaken about this, and we could also discover this by (e.g.) discovering a contradiction implied by one of our extensions. But this would be to discover either that the process which led to that belief did not reliably, context-independently, track necessary truth, and so was never an a priori warranting process, or it would be to discover that the fallible warrant conferred on the belief by the process was outweighed by the warrant of a more reliable process.[106] It would not, on the Platonist picture, be to discover that an a priori mathematical process was warrant-conferring in one circumstance, but not in another. Of course, Kitcher does not share the Platonist's view of mathematical ontology. But the apriorist who is also a Platonist is quite within his rights to say that an aspect of his position insulates that position against Kitcher's cases.

A further objection, one which follows on from the above response, is to grant that the Platonist is correct in saying that a process which reliably tracks the truth of necessary truths is one which cannot fail to satisfy (b). No process which actually is reliably warranting with respect to mathematical truths could become non-warranting. But this simply shifts the question to how we know which processes reliably track these truths. And this is the more interesting question, since if we don't know for any given process that it satisfies (b), then, as far as we know, that process is vulnerable to future (broad) experience. Our warrant for our belief that some process satisfied (b) surely is subject to change in the embedding context. If, for example, we come to believe that some axiom is wrong, then we will change some of our beliefs about processes which depend on that axiom. That is, we will change our belief about the warranting status of some process. Even if this would not be a case in which a warranting process becomes non-warranting, it would still be one in which a process was rejected as non-warranting. Thus, the actual practice of mathematics is vulnerable to future (broad) experience.

---

[106] More on the second possibility in Chapter Five.

The reliabilist can respond to this objection in two ways. The first way is by observing that the objection is based on a confusion between knowing and knowing that one knows. We have accepted from Kitcher a set of conditions on a priori warrant. Any process which satisfies these conditions is an a priori warranting process. It was no part of this account that we be able to tell which are the a priori warranting processes before they are to count as a priori. To say that the initial response just shifts the question is, from this perspective, to move the goalposts. A priori warranting processes are so, regardless of whether they are known to be so. Rational uncertainty over methodological issues, then, seems quite compatible with the claim that processes which reliably and context-independently track necessary truths can produce a priori warranted beliefs. To think otherwise is, in effect, to demand of such a process, before it is allowed to produce knowledge, that it be known to be such a process. It is, in short, to commit a restricted version of the KK fallacy.

The second response is to say that the objection insinuates that we have no idea which are the a priori processes, when in actual fact we do have good reason to believe (fallibly) of some processes that they satisfy (b), and thus are a priori warrants. We know from the fact that a sound proof logically entails the truth of the conclusion from the truth of the premisses, no matter what. We thus know that proof is an a priori process. Consider also the checks which would be made on a putative new rule of inference, or on a process which required an extension of an old concept. We can have confidence in the reliability of such a process on the basis of how it answers questions the answers to which we already know, how well its results fit with accepted mathematics, how well it survives attempts to show that it leads to contradiction, how it advances our mathematical knowledge, and so on. [107]

But, leaving aside the easy case of proof, and concentrating on methodological controversies, if we don't know which processes satisfy (b), isn't the way open for us to accept in-fact unreliable processes, and reject in fact reliable ones? Yes. This is part of the fallibilist picture. In all probability, there are reliable processes which we haven't discovered, or haven't been persuaded by. Possibly, some of the processes we now believe to be warranting are not (though this is less likely, given their fit with mathematics). So, it is quite consistent with Platonism and apriorism to admit that our beliefs about warranting processes can change. One who accepts each of reliabilism, Platonism, and apriorism just insists that this doesn't affect the warranting status of processes, only our confidence in them. For any a priori process, it only has to satisfy (b), we do not have to know that it does. Requiring that we know (b) to be satisfied would be relevant only to the question of whether we know that a given process is a priori warranting, a different question, and one for which we have good reason for confidence in our answer.

Examples of dispute over methodological procedures do represent cases of uncertainty about how mathematics as a whole is to proceed, but debate over methodological procedures does not, by itself represent a challenge to a Platonistic brand of apriorism. What is being debated is which processes we think are a priori warranting. Our debate cannot change the actual warranting status of a process, any

---

[107] In fact, in Chapter Five, I shall outline a rationalist account of the warranting process of intuition according to which these kinds of checks and balances are a priori.

more than our debate over the nature of a given rock formation can change a mountain into a molehill. To recap, if mathematical processes reliably track the truth of necessary truths, they are warranting. If, moreover, those processes do this without recourse to any element of narrowly experiential warrant, the warrant they confer is a priori. Thus, true beliefs produced by such processes constitute a priori mathematical knowledge. Nor, once we accept reliabilism, do we have to know that the prima facie a priori processes which produce our mathematical beliefs satisfy these conditions. It is enough if they in fact do so. Now, though I don't have a demonstrative argument that prima facie a priori mathematical processes are of this nature, there is good reason to think that they are. The logical tightness of sound proof, the fact that accepted mathematics is free from known contradiction, the usefulness of mathematics in the sciences, all of these give us reason to think that mathematical processes reliably track the necessary truths of mathematics.

Moreover, most of the warrant attaching to our beliefs about which processes are a priori warranting is itself (fallibly) a priori warranted. With the exception of the warrant emanating from the usefulness of mathematics in the physical sciences, the reasons we have for believing mathematical processes to be reliable are entirely intellectual reasons, taking nothing from narrow experience. How, in cases of such disagreement, can we know that the processes we demarcate as warranting are actually warranting? Well, as Kitcher (2000, 81) remarks, which path gets chosen will depend on how well it fits with the mathematics that has gone before. We might add to this that it will depend also on the relative importance of what has to be sacrificed against what is gained for each suggested route. Further, the choice will depend on the potential of each route for solving problems which have thus far gone unsolved, and on their respective fruitfulness for developing interesting new areas of mathematics. Here, as elsewhere, the reason for whatever confidence we have in our decision lies in the bet that the reason one option does all of this better than others is because it is true.

### 3.1.3

### The Argument From The Limitations Of Memory

Kitcher (1983, 44) has a further objection to proof apriorism , one which arises from reflection on one rationalist response to a problem like that presented by the Decay Picture. The uncertainty of our beliefs based on lengthy deductions had been a matter of concern for Descartes. His recommended solution was to continually go over such deductions until we develop the ability to see the entire proof in one single act of apprehension. Kitcher (1983, 45) correctly notes that there must be some upper limit on how much of this the human mind can do. For proofs which go beyond this limit, Descartes' suggestion will not help:

> As the reasoning develops, we are no longer able to keep in mind the evidence for the first principles; instead we have to "store" these principles, believing them now on the basis of the recollection that they were once established (1983, 44, original emphasis).

In other words, after a certain point, we believe some mathematical truths not on the basis of following a proof, but on the basis of remembering having followed one. Kitcher (1983, 45) continues:

> However this new process of recollection, although it normally warrants belief in the axiom, does not provide an a priori warrant for the belief. So, when we follow long proofs we lose our a priori warrant for their beginnings.

It is clear that, at least some of the time, in following a complicated proof, mathematicians cannot hold all the steps in mind at once. But, one might ask, so what? What reason is there to think that, once a cognizer must resort to memory in order to describe his warrant for a given belief, the warrant for that belief is therefore a posteriori? Suppose that the proof followed is sound, and each step has been correctly followed, and that the reasoner (correctly) remembers this. What justification is there for saying "Because memory was used, this does not count as a priori warrant?"

One can see why Kitcher would wish to deny memory the status of an a priori warrant. Memory does not satisfy condition (2b) on a priori warrant—given unkind experience, processes of the same type would not produce a warranted belief. If, for example, it seemed to you as though you remembered having followed a cogent proof of some result, but mathematicians everywhere countered that you couldn't be so remembering, as such a proof was impossible, or if your peers repeatedly told you that you were particularly vulnerable to mathematical flights of fancy, then for Kitcher, your belief in the conclusion of the proof would not be warranted. Since a process of the same type would not produce warrant under such circumstances, memory is not an a priori warrant.

Compare two cases where a mathematician correctly recalls having once followed a proof. In one case, his claim that he remembers is challenged, in the other, it is not. If these cases are to show that (2b) is not satisfied, they must show (i) that the same type of process was at work in both cases, and (ii) that memory confers warrant in the unchallenged case, but not in the challenged case. Again, these cases shows neither. The reasons for thinking that the same process is not at work in both cases are exactly as before. I shall spare the reader the repetition.

Unfortunately, the case for (ii) not being satisfied is not quite the same as it was earlier. What I would like to say is that, in the challenged case, if the proof was in fact sound, was in fact correctly followed and understood, and if these facts are correctly remembered, then we have still have a warranted belief. However, what enabled me to say this in response to Kitcher's case was the fact that sound proofs cannot possibly have defeaters, so one is still warranted in one's belief even if one does not consider apparent defeaters (since these could not be actual defeaters). This, as we all know, is not true of memory. Memory is prey to error in ways that following sound proofs is not. For example, a subject's memory might be unreliable in general, even though in a particular instance he correctly remembers having followed a proof. Knowing that this is a possibility, the subject whose memory-claim is challenged would not be warranted in stubbornly claiming that his memory had not led him astray.

There are, so far as I can see, three possible outcomes from a situation where one's memory of having once followed a proof is challenged. Firstly, one might, in response to the challenge, produce the proof (enough of it, at least, to indicate that one does, in fact, understand that proof). This outcome presents no problem for apriorism. Perhaps the only reason for thinking that it does is the thought that recourse to memory renders any justification empirical. But appeal to memory is a factor in all beliefs but those of immediate awareness (though how much of a factor it is in their <u>warrant</u>, as opposed to their production or recall, is controversial). If Kitcher's point is taken in this way, a priori warrant will be limited to a subset of propositions currently being attended to. This is not what is typically understood to be the issue between rationalists and empiricists. Kitcher might see this as grist for his empirical mill, but to view memory as automatically turning one's warrant into an empirical warrant is contentious. One may instead, with Burge, view memory as neutral between a priori and a posteriori warrant.[108] The function of memory, according to Burge, is to preserve content, which can be made available for reasoning (1993, 462). That memory is functioning correctly is not itself a premiss of the justification for any given belief. It is, like a correctly functioning brain, a necessary background condition for any warrant whatsoever. If there is no reason to think that, in a given case, memory is suspect, then the default position ought to be that it is functioning normally.[109]

Secondly, one may fail to produce a proof, but instead produce some other grounds for one's belief. For example, even if one could not produce a proof of the existence of irrational numbers, one could instead point to the fact that mathematicians have such proofs, or note that the result is recorded in many textbooks, and so on. In such a case, one would, in fact know that there were irrational numbers, but this knowledge would be empirically warranted.[110]

---

[108] See his "Content Preservation," Philosophical Review, V01. 102, No.4 (October 1993): 457-488.

[109] Burge (1993, 464-5) draws a distinction between two different types of role which memory can play. Where appeal to memory is in fact a premiss in a justification, this is "substantive" memory. Where memory is functioning solely to provide content for reasoning, but is not itself a part of the justification, the role is purely "preservative."

[110] Burge disputes this. For him, testimony is itself an a priori warrant. If, then, our warrant for belief in irrational numbers is made up of (i) testimony, from (ii) some source who does in fact know the proof a priori, then our belief is warranted a priori. I do not accept the claim that testimony provides a priori

Thirdly, there is the possibility that, when challenged, one may fail not only to produce a proof, but fail also to produce any other grounds for one's belief. In such cases, one's sole grounds for one's belief would be the memory of having seen a proof.[111] One's belief, then, would not count as knowing the result at all. But the claim that one does not know p a priori when one cannot produce any grounds whatsoever for p hardly presents a problem for apriorism. I may fail to know a priori that there are irrational numbers, but mathematicians do not. And this is enough for the apriorist. In short, no matter how we interpret the challenge allegedly posed to apriorism by appeal to memory, that challenge fails.

Summing up this section, then, we have seen that answering Kitcher's challenges to apriorism involves carefully distinguishing the origin of a belief from the warrant of that belief; it involves being clear about what is meant by experience, and experience-independence; it involves consideration of what role, if any, is played by memory in the warrant of beliefs other than those of immediate awareness; perhaps most importantly, it involves rejecting any notion of a priori warrant as indefeasible. Once we are clear about these points, Kitcher's cases pose no threat to apriorism.

---

warrant. There seems to me no good reason to attribute to myself a priori knowledge of some result in manifold theory, on the sole grounds that a friend in the mathematics department has told me that the result does in fact hold. Yet, if testimony provides a priori warrant, then, under these circumstances, I would know this result a priori—even though I may fail to understand it. If, however, the point is merely that testimony can transmit a priori knowledge, provided the receiver can understand it, then it is not testimony that is doing the warranting, but the intellectual process of understanding. Nor do I accept Burge's point that a belief may be warranted a priori, yet still be empirically defeasible. For me, insulation from narrow experience is one of the defining aspects of a priori warrant.

[111] Cases where one's sole ground for belief that p is the memory of having once seen a proof are different from another category of cases where one would not be able to produce a proof: cases of beliefs produced by the process of intuition. I will discuss how such beliefs are in fact warranted within a reliabilist epistemology in Chapter Five.

### 3.1.4

### "A Priori Knowledge Revisited" Examined

Kitcher in his more recent "A Priori Knowledge Revisited," has honed his attacks on apriorism about mathematical proof. The main strategy pursued there is to distinguish between the "Strong Conception" of the a priori, which accepts all of Kitcher's conditions, and the "Weak Conception," which accepts only condition (a), and to argue that, the Weak Conception runs into various problems. What would rescue the Weak Conception, Kitcher argues, would be the inclusion of (b), but this, Kitcher contends, would ultimately mean adopting the Strong Conception, and the Strong Conception, Kitcher believes, does not count mathematical knowledge as a priori. If what I have outlined above is broadly right, there is third conception available, let us call it the Moderate Conception, which accepts conditions (a) and (b), but not (c).

Kitcher's argument that no such moderate conception exists goes as follows. If there were a process which satisfied both (a) and (b), but not (c), this would mean that one could have experiences sufficient to believe that p, one would be warranted in so believing p, but that p was false. Kitcher calls any such set of experiences "e"; he (2000, 72) then asks us to consider the following extension of e, e*:

> [T]he next stage of one's life consists in an encounter with an oracle who demonstrates power to answer vast numbers of significant questions, who testifies to the falsity of p, and who offers whatever can be done to show directly that p is false. e* is sufficient for p, so, by assumption, the process is available given e* and its warranting power is unaffected. But, in order to believe that p, one must override extremely strong evidence to the contrary. How can the process give one a license to override when one would have been epistemically better off not overriding?

Kitcher (2000, 72) goes on to claim that the person who overrides the performance of the oracle is being dogmatic in insisting on his belief in p, and that the belief is no longer warranted. The case, he concludes shows that a process cannot satisfy (b) unless it also satisfies (c).

From the point of view of one who holds that a priori warrant is fallible, and who accepts (b) (when read as requiring reliability of a priori warrants), there are two responses to this case. The first response focuses on the fact that the case is one where there are conflicting warrants. Clause (b), as we have seen, demands that an a priori warrant would warrant p. This, however, is not to say that an a priori warranting process will, always and forever, provide sufficient warrant for belief in p. In particular, such a process (fallibly) warrants p until such a time as a warrant for belief in not-p arises. When a not-p warrant is made available, the question of which is the stronger warrant arises. In Kitcher's case, the fallibilist apriorist can reply that belief in p is warranted a priori in e. In the extension, e*, p is warranted (with no reason for doubt) until the subject meets the oracle, at which point, he is presented with warrant for not-p. At this point there are two cases. The oracle might just tell the subject that p is false, without giving a reason for this claim. If our subject is armed with something like a proof (since p is supposed to be false, it cannot of course be a sound proof), then it remains to be shown that the oracle's testimony outweighs the misleading proof. If on the other hand, the oracle presents a refutation of p, and our subject can understand this refutation, and see that it is sound, our subject's warrant for p is undermined. Thus, the initial process does offer some degree of warrant, but this warrant is defeated by the, ex hypothesi, stronger warrant for not-p offered by the oracle. Thus the case is not, in fact, one where the same process warrants a belief in one set of circumstances, but not in another. The process does confer warrant on p, but that warrant is outweighed by another, more reliable process. Since the case is not one where a warranting process occurs without conferring warrant on p, it is not one where (b) is violated. Kitcher's reason for thinking that it is lies in his also thinking that the warrant conferred by an a priori warranting process remains the same, come what may. The fallibilist apriorist holds rather, that a process may indeed confer warrant, but yet have that warrant undermined by other conflicting warrants.

The second response focuses on the subject's epistemic situation after the oracle has given him reason for not-p which is strong enough to defeat his warrant for p. If the subject continues to believe p at this point, his belief is, on balance, unwarranted. But note that the process which sustains belief in p at this point is not simply the a priori process which he earlier thought to warrant p. For now, in addition to this, he must either ignore, or find reasons to reject, the oracle's refutation of p. And, in either case, belief in p is sustained by a <u>different process</u> in e* than in e. Thus Kitcher has not given us an instance of the same a priori warranting process sustaining warranted belief in p in one set of circumstances, while failing to do so in others. He has, then, failed to provide a case the solution of which makes (c) mandatory.

I take the above to show that (c), as it is not a consequence of Kitcher's other conditions on a priori warrant, may be rejected, even by one who accepts (b). Kitcher's <u>de facto</u> argument that there is no moderate conception thus fails. My strategy in what follows will be to show that the criticisms Kitcher aims at the Weak Conception of the a priori fail when construed as objections against the Moderate Conception, and that, even if mathematical knowledge does not count as a priori on the Strong Conception, this is not enough to establish that mathematical knowledge is not a priori.[112]

### 3.1.4.1

### Experience-Independence and the Moderate Conception of the A Priori

Kitcher (2000, 77) argues that the Weak Conception of the a priori abandons experience-independence. This is because giving up (b) abandons the idea that what is a priori is not vulnerable to future experience. Since the idea that the a priori is experience-independent is a central part of the traditional notion of the a priori, the Weak Conception misses out on a central facet of a priori knowledge, which Kitcher's Strong Conception succeeds in capturing. It should by now be clear that the Moderate Conception does not abandon experience-independence in any sense which would be damaging to the apriorist. Nothing can genuinely show an a priori warranting mathematical process to be non-warranting. So, the warranting powers of such processes are not experience-dependent in any sense. This suffices to satisfy (b).

But a priori warranting processes produce beliefs, and those beliefs are open to being undermined, so won't our a priori warrant be hostage to future experience after all? The response has already been given. A priori warranted beliefs are not open to being defeated by narrow experiential reasons, and this is the only sense of 'experience'

[112] I will discuss here only two of the three arguments Kitcher offers against the Weak Conception. The third argument is to the effect that the Weak Conception cannot correctly account for certain kinds of probability reasoning. It does not advance the main point of the paper: that mathematical knowledge either counts as a priori, but tradition-dependent, on the Weak Conception, or fails to be a priori, on the Strong Conception. Since I am arguing that the Moderate Conception avoids both of these conclusions, and thus counts mathematical knowledge as both a priori and tradition-independent, and since the probability problem arises only for the Weak Conception, I need not address this argument.

that counts in the debate between apriorists and their opponents. A priori warranted beliefs are open to being defeated by intellectual considerations—for reasons of broad experience, that is. This is part of the fallibilist conception of the a priori. But far from being a damaging concession, this is a good thing. It allows mathematicians to revise their beliefs in the face of "experiences" such as being presented with a refutation of one of their beliefs. The Moderate Conception maintains narrow experience-independence for a priori warranted beliefs, and thus does not allow a priori warranted beliefs to be undermined by narrow experience. This, I have urged is the only kind of experience relevant to the debate between apriorists and their opponents.

#### 3.1.4.2

#### Tradition-dependence

The Weak Conception, Kitcher argues, fails to preserve an important element of traditional thought concerning a priori warrant. To be more specific, the Weak Conception does not capture the notion of the tradition-independence of a priori mathematical warrant. Kitcher (2000, 81) explains tradition-independence thus:

> [A] person's knowledge is independent of socio-historical tradition just in case that person could have had the knowledge, even given socialization in a different tradition, provided only that the socialization made it possible to entertain the propositions known.

Kitcher observes that Frege, along with subsequent writers on the foundations of mathematics, was attempting to provide warrant for our mathematical beliefs which was independent of tradition. In order to coherently capture the idea of the tradition-independence of mathematical warrant, and to make sense of the Fregean project, the apriorist needs to insist on (b), since any processes satisfying (b) "would provide warrants for belief that do not depend in any way on the particularities of what X has absorbed from the past" (Kitcher, 2000, 81).

But this is exactly what the Weak Conception gives up. The Weak Conception allows that there are possible experiences in which processes which would normally warrant mathematical beliefs fail to do so. That the Weak Conception allows this is shown by the following case. Imagine that the generation of mathematicians preceding our own had some such experience. Either, argues Kitcher, they tell us that the processes they have (experientially) found reason to mistrust are unreliable, or they do not. In either case, we would not be warranted in using these processes. If they tell us they are unreliable, and we use them anyway, we are clearly unwarranted. If they do not tell us, and we use the processes, then we are still unwarranted, as we are being raised to use processes which are both unreliable and known to be so (though not by us). Either way, we are unwarranted in the hypothetical case, in using processes which, in the actual world, are warranted. Thus, Kitcher concludes, the Weak Conception of a priori knowledge has the consequence that mathematical knowledge is tradition-dependent. Worse still for the apriorist, Kitcher further holds that this tradition-dependence of mathematical knowledge is of a sort which connects with experience in two distinct ways. Clearly, the apriorist is in trouble if mathematics is tradition-dependent in any such sense, though, as we shall see, not all types of tradition-dependence need trouble him.

It should already be clear that since the Moderate Conception includes (b), it will not be vulnerable to Kitcher's argument. As Kitcher notes, an account of a priori warrant which requires of an a priori warranting process that if it were to produce the belief that p, then it would warrant p, would not allow a process to count as a priori warranting if it could warrant belief in some circumstances, but not in others. So, since a priori mathematical processes satisfy (b), our ancestors could not have had experiences which showed that an a priori warranting process was unreliable. Thus, mathematical warrant, on a conception of the a priori which includes (b), is tradition-independent.

Kitcher, of course, is not out to argue that the inclusion of (b) saves the day for the mathematical apriorist. His position is not merely that the Weak Conception yields the conclusion that mathematical knowledge is tradition-dependent. Rather, the position is that mathematical knowledge actually is tradition-dependent. Kitcher also holds that mathematical knowledge fails to satisfy condition (b) (if it were to satisfy condition (b), it would be tradition-independent). What reason then, does he offer for thinking that a conception which contains (b), such as the Moderate Conception, will yield either the conclusion that mathematics is tradition-dependent, or that mathematical knowledge fails (b)?

I see only two possible lines of argument.[113] The first line of argument is that we know that mathematical knowledge does not satisfy condition (b) because Kitcher (1983) has offered cases where experience undermines the warrant attached to mathematical beliefs. But we have argued that these cases do not show mathematical

[113] Kitcher (personal communication) has confirmed that these are the lines of argument he had in mind.

knowledge fails (b).[114] This line of argument in support of the claim that mathematical knowledge fails (b) is thus unsupported.

The second line of argument is to appeal directly to the plausibility of the tradition-dependence of mathematical knowledge. If mathematical knowledge were to be genuinely tradition-dependent in Kitcher's sense, this would mean that there are elements of mathematical knowledge which would not be knowable by mathematicians had they been socialized in a way different from how they actually were socialized. A plausible case that mathematical knowledge is tradition-dependent in this way would ipso facto be a plausible case that mathematical knowledge depends on experience, and thus that it fails to satisfy (b).

What reason is there to hold that mathematical knowledge might be tradition-dependent in this sense? Kitcher holds that a consequence of tradition-dependence is that the warranting power of the processes used by mathematicians is tradition-dependent, and that this feature of tradition-dependence connects mathematical knowledge with experience. Either of these claims, if true, would mean that mathematical knowledge did not satisfy (b). Is mathematics tradition-dependent in a sense that has either of these consequences? I shall examine the two claims in order.

Firstly, then, is the warranting power of mathematical processes tradition-dependent? Kitcher (2000, 84) claims that "our knowledge is dependent on the tradition in which we stand, and … our [mathematical] knowledge … depends on the experiences of historical figures who have played important roles in past inquiry." He

[114] Strictly speaking, we have so far only argued this for non-basic mathematical processes. We will take up the arguments concerning basic mathematical processes in Chapter Five.

(2000, 84) further claims that "the warranting power of the processes of thought they [i.e., apriorists] take to underlie mathematical knowledge depends on the experiences of those who came before us in the mathematical tradition" (emphasis in original). If the warranting power of the processes used by mathematicians depends on the experiences of those who came before them, and if those experiences might have been different, then the warranting power of those processes would be tradition-dependent.

But what experiences of historical figures are in question here? The "experiences" of discovering that some processes are a priori warranting, and the "experiences" of using these processes to discover new results, which are thus justified a priori? The "experiences" of discovering that some processes, thought to be warranting, either are not genuinely warranting, or else do not confer warrant outside of certain restricted areas of application? The "experiences" of thinking through concepts, and discovering that certain things follow from others, or lead to desirable results, or to paradox, or other such "experiences?" To the extent that these are experiences at all, they are all kinds of what I have been calling broad experience. The findings involved in such experiences were discovered using methods that seem clearly a priori.

Moreover, such "experiences" do not show that the warrant-conferring power of a priori processes is tradition-dependent, only that our thought about such processes is. That is, thanks to the efforts of earlier mathematicians, we now know a great deal about which processes are warranting, and which are not. This knowledge is tradition-dependent, but this is a harmless admission. The warranting power of the processes themselves is not.

For sure, it might have been that a historical figure hadn't made a given discovery, and that no-one else made it instead. But this does not show that the knowledge in question is tradition-dependent in Kitcher's sense, since this does not preclude the possibility of a mathematician coming to discover, and thus know, the finding in question, even given different socialization. Any given item of a priori warranted knowledge was produced by a process which, I have argued, is, and cannot fail to be, warrant-conferring. It would be possible for a mathematician, even given different socialization, to use any such process, once that socialization enabled him to entertain whatever propositions were needed to use the process in question. And, by the same argument from Platonist ontology given above, together with the fact that it is not required that one know of a warranting process that it is warranting, beliefs produced by this process would be warranted—even if they were at odds with the beliefs of his tradition. None of these kinds of experiences, then, show that the warranting powers of prima facie a priori mathematical processes are tradition-dependent.

Coming at things from a slightly different angle, Kitcher (2000, 76) claims that the entire conceptual framework of mathematics is warranted experientially: "we might reasonably view our—allegedly a priori—knowledge as dependent on experience, the experience that warranted adoption of the conceptual framework on which we now draw." But this route fares no better. The conceptual framework we have adopted has been shown, by the a priori investigations of those who came before, to be warranted. The processes which warrant the adoption of the framework are intuitively a priori, and there is no reason to think either that a mathematician, given different socialization, and the ability to entertain the relevant propositions could not have come to know them, or that they do not satisfy (b). In fact, this is supported by what Kitcher (2000, 76) thinks is the most plausible way of looking at mathematical innovation: "[mathematical innovators] are warranted in proposing new concepts and principles through an often lengthy process of demonstrating that their new ideas play a fruitful role within inquiry." This method of conferring warrant on mathematical beliefs is entirely a priori. The proposed principle or process is examined to see that it is free from contradiction, and that it is fruitful. This is a matter of seeing what follows from it, and what support can be offered for it from generally accepted mathematics. This method is warranting, though fallible. On Platonist assumptions, if it is indeed warranting (as there is good reason to think) it cannot become non-warranting. Thus, it satisfies (b), has its warranting power independently of tradition, and counts as moderate-conception a priori.

Mathematical knowledge is tradition-dependent in at least two ways, but it can be shown that these are harmless ways. The first way has to do with the accumulation of that knowledge. For any given item of knowledge, if someone hadn't discovered it, then, unless the current generation of mathematicians themselves were to discover it, they would not now know it. The caveat here shows that this sense is harmless, for in order to show that any given item of knowledge would not be know<u>able</u>, it would first have to be shown that the item in question would not be <u>discoverable</u> by a present-day mathematician. In order to show this, it would have to be shown that a present-day mathematician could not use the a priori warranting process which actually led to the discovery of the item in question. For, given that it is warranting, once he used it, and it produced the relevant belief, he would thereby know the item. If the different socialization envisaged allows the mathematician to have whatever concepts are needed to use the process, it is going to be extremely difficult to argue that he could not use the process.

The second way in which mathematical knowledge is dependent on tradition is that mathematicians must undergo training in their discipline, and must demonstrate that they have learned a sufficient amount of what the tradition has bequeathed. This too is a harmless sense, because what's being passed on is the use of a priori warranting processes, and the results these processes have warranted. Perhaps some number of the processes that get passed on actually do not reliably track the truth of necessary propositions, and thus, fail to be a priori warranting processes. But for those processes which are genuinely a priori warranting processes, they cannot fail condition (b), and thus cannot depend for their warranting powers on the tradition. Moreover, it would be possible for a mathematician, even given different socialization, once it enabled him to obtain the concepts needed to use one such process, to use that process. And once again, by the same argument from Platonist ontology given above, beliefs produced by this process would be warranted—even if they were at odds with the beliefs of his tradition.

What about Kitcher's (2000, 83-4) claim that tradition-dependence links mathematical knowledge to experience in two senses? Even if the warranting power of mathematical processes were not tradition-dependent, perhaps there is some other way in which (an otherwise acceptable sense of) tradition-dependence connected mathematical warrant with experience. If this were so, it would obviously create a problem for the apriorist.

The first way mathematics depends on experience, according to Kitcher (2000, 84), is that the "ultimate starting points [for mathematical beliefs] lie in those scattered perceptions that began the whole show." It may well be true that something like this explains the genesis of our mathematical beliefs. But, the apriorist holds that this has no bearing on the further question of whether there are a priori warranting processes. The fact that one can have empirical warrant for some mathematical belief (e.g., inductive warrant for the claim that 2+2=4) simply has no bearing on whether that belief is of a kind which can be warranted a priori. Likewise, the fact that one needs some experience in order to come by the concepts required for mathematical thinking has no bearing on the warrant that may or may not attach to beliefs involving those concepts.

The second way in which Kitcher sees mathematical warrant as depending on experience depends on society's division of labor. We have, he (2000, 84) tells us, learned from experience that having a specialized group of people who extend and articulate mathematical languages is a good thing, since doing things in this way promotes inquiry. But this does not show that mathematical knowledge depends on experience. It shows that experience warrants our belief that having specialized mathematicians together in the same geographical area leads to desirable developments in mathematical thought. Experience shows that this way of doing things leads to a desirable rate of development of mathematical thought, one which is probably better than many alternative ways of doing things. But experience does not tell us anything about how the results themselves are warranted. And this is what would need to be shown for mathematics to be dependent on experience in any way that would threaten its a priori status.

There is one last point to consider. After Kitcher makes his claim that mathematical warrant is tradition-dependent, he (2000, 84) writes:

> Once we see this, we'll recognize that the issue isn't one of apriorism versus empiricism, but of apriorism versus historicism, and here the interesting question is whether one can find, for logic, mathematics, or anything else, some tradition-independent warrant, something that will meet the requirements that Descartes and Frege hoped to satisfy, in short, something that will answer to the Strong Conception.

What is the apriorist to make of this challenge? Does he need to find something that answers to the Strong Conception? Once we give up (c), and embrace a brand of fallibilism, we do not need to. The search for certain foundations, which are known to be so, is not one the fallibilist need undertake. Moreover, if one is reliabilist into the bargain, there is a fairly easy route to arguing that the absence of such certainty does not prevent us from having a priori mathematical knowledge. If intuition is reliable, and if proofs are reliable, then we have a priori knowledge. Whether or not any one person can display the structure of that warrant, a la Frege, is relevant only to the question of whether we know that we have a priori knowledge. Although the success of mathematics, together with the logical tightness of mathematical proof, would tend to support a positive answer to the question of whether we have reason to believe our extant mathematical processes to be reliable.

Traditional thinking on the a priori is a "complicated mess," according to Kitcher (2000, 83 n24), and it makes little difference whether we say that mathematics is not a priori on the Strong Conception, or argue that it is a priori on the Weak Conception, so long as we admit its tradition-dependence (2000, 85). Kitcher may well be correct in his diagnosis of traditional thought on the a priori. Certainly, he has done much to force the apriorist to sharpen his thought. It seems to me, in fact, that what Kitcher has done is to show that the choice is that between some form of mathematical empiricism, and a Platonist brand of apriorism. Consider how heavily I have leant on the Platonist ontology of mathematics to rebut Kitcher's objections to apriorism. This lends weight to the idea that if one wants a version of mathematical apriorism according to which mathematical knowledge is tradition-independent (in Kitcher's sense), and mind-independent, then one must embrace Platonism.

Regardless of this last point, however, the dilemma that Kitcher offers between mathematical knowledge being a priori in the weak sense, but tradition-dependent, or simply failing to be a priori on the strong sense, is false. The Moderate Conception preserves the tradition-independence of a priori warrant, while still counting mathematical warrant as a priori.

## 3.2

## Human Limitations and Computers: Tymoczko and The Four-Colour Theorem

The second source of argument against the apriority of mathematical proof to be discussed here is Tymoczko (1979). There, Thomas Tymoczko presents two (closely related) objections to the apriority of proof, both arising out of the introduction into mathematics of computer-assisted proofs. Firstly, there is the claim that the unsurveyability of some computer proofs, such as the one used to prove the four-colour problem, shows that such proofs must be established on empirical grounds, and thus cannot be like traditional, a priori proofs. Secondly, there is the claim that in appealing to computers, one is appealing to a physical experiment, and so proofs which require such an appeal are not established a priori. In order to bolster this second claim, Tymoczko offers three supporting arguments. The first two of these are that the possibility of (a) hardware and (b) programming errors means that a proof which appeals to computer must be a posteriori. Thirdly he argues that since the truth of the four-colour theorem (4CT) is not as certain as the truth of many empirical propositions, it cannot be a priori. I shall discuss each of Tymoczko's arguments in turn.

### 3.2.1

### The Argument From Unsurveyability

The problem to which the 4CT provided the solution was the problem of whether all countries on a map could be coloured with only four colours in such a way that no neighbouring countries were of the same colour. The positive answer to this question, accepted since the late seventies, was provided, in part, with the aid of a computer. Consideration of this aspect of the 4CT leads Tymoczko (1979, 69) to the following facts:

1. Proofs, in the traditional, a priori, sense are surveyable.
2. No mathematician has ever seen, or could ever see, a proof of the 4CT (the 4CT is not surveyable, in the traditional sense).
3. Because the proof of the 4CT is not surveyable, the use of a computer is indispensable.
4. The 4CT leads to a genuine extension of mathematical knowledge. It is not mere calculation, but a proof of a substantial new result.

Not only is it the case, then, that the 4CT has not been surveyed by any human mathematician, no human mathematician ever could survey it. No computer has printed out the complete proof. Even if were to be printed out, the print-out, because of its sheer length, wouldn't be of any use to human mathematicians. It cannot be checked step by step like a traditional proof. In effect, Tymoczko's argument from unsurveyability identifies surveyability as a crucial feature of traditional, a priori proofs, notes that the proof of the 4CT cannot be surveyed, and concludes that the 4CT is not a traditional, a priori proof. The truth of the 4CT is knowable only a posteriori (Tymoczko, 1979, 77).

It is certainly true that no human mathematician could survey the entire proof. But why should this be supposed to cut so much epistemological ice? Could we not see the appeal to computer as a convenient abbreviation in the proof—akin to reference to results published elsewhere, or to a traditional appeal to a log-table? Tymoczko (1979, 70) answers this question negatively, supporting his verdict by noting that, in traditional proofs, any such abbreviation is itself backed by a surveyable proof. Furthermore, "[i]n principle this surveyable backing is available to any member of the mathematical community" (Tymoczko, 1979, 70, emphasis added). In contrast, that portion of the proof of the 4CT which is carried out by computer is not surveyable. It is not recorded in the archives, and is anyway too long for any one mathematician to read.

This unsurveyability may be circumvented in one of at least two ways. Firstly, although the entire proof cannot be surveyed, any given part of it could be. The proof may be sampled at a number of points, and the reasoning checked. Secondly, we could print out the proof and give a small part of it to each member of the community of mathematicians. In this way, the proof would be surveyable, not by one human mathematicians, but by all of them acting together.

One possible objection to these methods of surveying the 4CT is that although implementing such methods would indeed result in the proof being surveyed, there still would not be any one human mathematician who had surveyed the entire proof. Thus, the 4CT still lacks a feature possessed by traditional, a priori proofs—including those using traditional shorthand conveniences. To see that this objection fails, consider again the "in principle" surveyability which, according to Tymoczko attaches to abbreviated traditional proofs. How is the "in principle" surveyability to which he appeals to be explicated in such a way that a mathematician can, in principle, survey an extremely long traditional proof, but cannot survey the 4CT? This point is especially pressing if the proof in question is one that relies on other complex proofs that have gone before. The answer cannot simply be that a mathematician could, in the space of a lifetime, survey the traditional proof, but not the 4CT. There might, after all, be proofs which, given their dependence on other long and complex proofs, could not be traced back to axioms or self-evident principles in a human lifetime. As Margarita Levin[115] remarks in her response to Tymoczko, the whole of traditional mathematics is very probably unsurveyable by any one human mathematician. This does not render a posteriori the proofs of traditional mathematics. [116]

[115] Margarita Levin, "On Tymoczko's Argument For Mathematical Empiricism, Philosophical Studies 39 (1981): 84.

[116] Margarita Levin notes that the fact that the whole of traditional mathematics might be unsurveyable by a single human mathematician could itself be used to mount an argument for mathematical empiricism. However, as will become clear, when a priori truth is understood along the lines offered here, this is no longer a possibility.

But if "in principle" is not to be explicated in this way, there is no obvious way to explicate it such that traditional proofs are "in principle" surveyable, but the 4CT is not. And, absent such an explication, Tymoczko has no objection to the suggestions made above as to how to circumvent the 4CT's unsurveyability. Granted this, we can reinstate the response that the appeal to computer is a shorthand convenience in the proof, akin to reference to results published elsewhere. If traditional proofs are knowable a priori because they are "in principle," even though not in practice, surveyable, what is wrong with seeing the 4CT as knowable a priori? It too is "in principle," if not in practice, surveyable.

In an attempt to support his claim that unsurveyable computer results must be established on empirical grounds, and thus are not knowable a priori, Tymoczko (1979, 71) offers the following analogy. We are asked to consider the hypothetical history of Martian mathematics. This history proceeded much as our own, until the arrival on Mars of a genius named Simon. Under Simon's guidance, Martian mathematics made great strides. Eventually, however, Simon began producing proofs which included phrases such as "Proof is too long to include here, but I have verified it myself." On occasion, Martian mathematicians could not reconstruct the reasoning which had gone on behind these new "lemmas," but such was Simon's prestige that they accepted these results under the general heading "Simon Says." Tymoczko concludes firstly that Martian mathematics under Simon is not a legitimate development of mathematics, and secondly, that we are in a similar position with respect to our appeal to unsurveyable computer proofs.

Given Tymoczko's description of the case, the first of his conclusions seems unobjectionable. The second, however, is simply false. There are important differences between the position we are in with regard to computers, and the one that the Martians are in with regard to Simon. One such difference is that we (i.e. members of our community) designed the computer. The fact that we designed the computer gives rise to other, more significant differences between the computer case and the Simon case. The computer is designed to do a certain job: it is designed to follow programs. We know what program the computer is supposed to be executing. We can, therefore, as often as we choose, check to ensure that the computer is in fact executing this program. Furthermore, the proofs a given computer discovers can be checked by others. Random portions of these proofs can be hand-checked by human mathematicians. The program may be tested by having a computer verify already established proofs. If the computer failed such tests, we would no longer believe its results. In the case of Simon, the Martian mathematicians did not design him, and at least some of the tests available in the case of computers are not available in his case. They cannot, for example, open him up and examine his inner workings. Nor can they assume, from the fact that he has given the correct answers in verifiable cases, that his answers in the unverifiable cases are likewise correct. He may be lying. He may have lost his ability, and be attempting to protect his prestige. It is true that, in the case of the computer, it might somehow have lost its ability, but this is exactly the sort of thing that can be checked, if desired, simply by running tests on whether it is executing its program. No such check is available in the case of Simon, since no-one knows what his "program" might be.

What divides our case from that of the Martians is that we know that computers are in-principle reasoners. And here, the 'in-principle' serves simply to indicate that the computer is not subject to the same speed and mortality limitations as are human mathematicians.

Although Tymoczko (1979, 74) notes some of these features of computer proofs, he doesn't seem to see them as damaging to his analogy. This, given that the whole point of introducing the Simon analogy was to support the claim that unsurveyable proofs are not a priori, is rather odd. For Tymoczko himself (1979, 72) observes that computers are, unlike Simon, not just authority, but warranted authority. It seems strange to suppose that comparing an appeal to a warranted authority with an appeal to an unwarranted authority could tell us anything about whether the warrant attaching to the warranted authority was a priori or a posteriori. The cases are too dissimilar for the analogy to be particularly enlightening. All they have in common is that they are both unsurveyable. Whether this unsurveyability in the case of the 4CT renders it a posteriori depends on the methods used in the proof. Comparing the 4CT to an unsurveyable and unwarranted method sheds no light on this.

If, as has been argued, the unsurveyability of the 4CT does not leave it lying outside of the set of traditional a priori proofs, we are left with something of a dilemma. We either have to say that the 4CT is known a priori, just as we say that proofs with other shorthand conveniences are known a priori. This move I find unappealing. The alternative is to say, along with Tymoczko (1979, 77), that a priori truths are those truths which can be known a priori, and then to argue that the 4CT is just such an a priori truth. This I will attempt in the following section.

### 3.2.2

### Does the Appeal to Computer Make the 4CT an Experimental Result?

There is then, no reason to think that the unsurveyability of the 4CT presents an insurmountable problem for mathematical apriorism. But what of Tymoczko's (1979, 63) claim that an appeal to computers is an empirical premiss, a report on an experiment? Tymoczko (1979, 75) remarks that what makes it reasonable to accept the 4CT is that although there is no surveyable proof, we know that there is a formal proof of it. This knowledge is acquired through the use of computers. Tymoczko's case that the appeal to computer is ultimately a report on a successful experiment rests on his further claims that the 4CT involves an appeal to the reliability of computers, and that this latter appeal is an empirical matter.

Of course, the mere fact that a computer "experiment" has yielded a genuine extension of mathematical knowledge is not in itself a threat to apriorism. A threat would be posed only if (a) the "experiment" were an experiment in something like the sense in which the term is used in the physical sciences, and (b) the proof were not discoverable using purely a priori processes. Mathematical apriorism, after all, is not committed to the claim that mathematical truths cannot possibly be known a posteriori; one could come to know a mathematical truth through being told by a reliable expert that it holds. Knowledge thus transmitted would be a posteriori. Mathematical apriorism is committed only to the claim that such truths <u>can</u> be known a priori. Tymoczko's case faces two questions, then. Is the 4CT accurately describable as an experiment, in any sense that would mean that it was a posteriori, and does the discovery of the proof involve methods which are not a priori?

Tymoczko himself (1979, 78) notes one strange consequence of seeing the 4CT along the lines of an experiment. Having claimed that the 4CT blurs the distinction between an alleged a priori discipline like mathematics, and natural science, Tymoczko (1979, 78) runs up against the problem of what to say about experiments in pure mathematics:

> It is easy to see how experiments play a role in the arguments of physical theory. The physical theory can predict phenomena of space-time which equipment can be designed to register. Are we to say that the computer registered a phenomenon of mathematical space? If not, then how else are we to explain the role of experiment in mathematics?

Few philosophers would wish to cleave to the first horn of Tymoczko's dilemma. He himself does not, but nor does he have any suggestions as to how to go about unpacking the second horn. Instead, he (1979, 78) decides to "simply note these puzzles as among the consequences of the 4CT."

But the fact is that these puzzles are among the consequences of the 4CT only if that theorem really is accurately describable as an experiment. If we do not classify it as such, the puzzles disappear. And surely having to affirm the disjunction that we have either registered phenomena in mathematical space, or else really have no idea what we mean when we talk about an experiment in mathematics, ought to count against classifying the 4CT as an experiment.

Moreover, there are further disanalogies between physical experiments and the use of computers in proving the 4CT. A rough characterization of what an experiment is in the physical sciences would be that it is an attempt to control and isolate part of nature, in order to learn, through observation, about that part of nature, or about some other part of nature to which the controlled and isolated part is causally connected. That is, an experiment, in the sciences in which the notion has a reasonably well-understood use, involves all of the following: the control and isolation of nature, observation, prediction, and the correlation of causes and effects. None of these features are present in the appeal to computers in the 4CT.

Take the correlation of cause and effect. Mathematics has no place for these notions. Whether mathematics is ultimately about abstract objects, ideal mental constructions, or fictional objects, mathematical objects neither are nor have causes or effects. Whatever factors are involved in the 4CT's appeal to computers, then, the correlation of causes and effects is surely not amongst them.

Things are only slightly different when it comes to observation. While there is a role for observation in the 4CT, this role does not involve the sort of observation that is central to the notion of experiment. In an experiment, the purpose of observation is to learn either about the phenomena being observed, or about phenomena causally connected with the observed phenomena. For the reasons given in the previous paragraph, the role of observation in the 4CT cannot be of the second sort. But nor can it be of the first sort. The role of observation in the 4CT is simply to read the output provided by the computer, a result which the computer has arrived at by some other non-observational method. That output is the computer's report that it has established a certain truth about mathematical reality. That is, it is a report on a concluded investigation. But this is not the primary role of observation in the physical sciences. There, observation is not merely observation of a report that a truth has or has not been established via some other method, it is itself the method of establishing truth or falsity. Given this vital difference between how observation is used in the physical sciences, and how it is used in the proof of the 4CT, it seems fair to say that observation, in the relevant sense, plays no role in the proof of the 4CT. Indeed, if the presence of observation, in this sense, were a sufficient condition for inquiry to count as experiment, then any "observation" that any result in mathematics had been established would count as an experiment, and there would be no need to appeal to the role of the computer in the 4CT to undermine mathematical apriorism. This, surely, is an unacceptable consequence of classifying the 4CT as an experiment.

Nor is there a role for prediction in the 4CT. Scientific hypotheses make predictions, both about what will happen when a body is acted on by a force, and about what an observer might see under certain conditions. Thus, in astronomy, one might find an experimental claim along the lines of "if the telescope is rotated through 12 degrees, body X becomes visible." Mathematical hypotheses, on the other hand, make no observational predictions.[117]

Again, while it is true that part of nature—the computer—is being controlled while the computer runs through the proof of the 4CT, it is not true in the relevant sense. Nature is controlled in the physical sciences in order to learn something: either about something that is causally connected with that part of nature, or about that part of

[117] Thanks to Michael Levin for pointing out my omission of the role of prediction in the physical sciences in a previous draft of this chapter.

nature itself. Neither of these motivations is at work in the 4CT. There, the controlling is not undertaken in order to learn about the computer itself, nor about some part of nature to which it is causally connected. Nor even is it to prevent other parts of nature from interfering with the relationship being studied. Rather, the computer is being controlled in such a way so as to model mathematical and logical relations, holding between premisses and conclusions. It is these latter that are important to the proof of the 4CT, not the physical part of nature being controlled. The hardware of the computer is irrelevant. Provided only that it is capable of running a program which instantiates the relevant mathematical and logical relations, any hardware—i.e., any part of nature—would do. This is a striking difference from the control of nature practiced in physical experiments. There, the part of nature to be controlled will be determined by the physical relations one wishes to study.

Might one hope to save something of Tymoczko's claim by putting another interpretation on the role of the computer in the 4CT? Instead of holding that the appeal to computer constitutes an experiment, might one instead see the computer as an instrument of observation, somewhat similar to a telescope?[118] This interpretation fails, however, because the analogy is a weak one. An instrument of observation, such as a telescope, fulfils its function by virtue of being causally linked in some way to the phenomena observed. The telescope allows us to detect light rays which emanate from the observed body by virtue of the fact that it interacts causally with physical entities, themselves causally connected to the observed phenomenon. Since mathematical

[118] Thanks to Michael Levin, who raised this possibility (without endorsing it, I hasten to add).

phenomena are neither physical nor causally efficacious, the analogy between the role of the computer in the 4CT and the telescope breaks down. A successful analogy here would have to show the role of the computer in the 4CT to be similar to an experimental instrument of observation which did not require a causal link of some sort between the phenomenon and the instrument. It is difficult even to imagine how this could be made coherent.

Given these striking dissimilarities between physical experiments and the 4CT—dissimilarities, moreover, having to do with features central to the very notion of what an experiment is—it seems at best misleading to think of the 4CT in terms of an experiment. Tymoczko himself (1979, 77) admits that it is strange to call the 4CT's appeal to computer an experiment, but yet he uses the claim that it is an experiment in order to argue that the 4CT is not an a priori truth. But conversely, if a plausible explication of a priori truth is available which counts the 4CT as a priori, this would tend further to support the claim that there is no "experiment" involved in establishing the truth of the 4CT.[119]

And there is a conception of the a priori which does not classify the 4CT as a posteriori, and hence avoids all of the above problems. An a priori truth, as Tymoczko (1979, 77) notes, is one which can be known without recourse to experiential premisses.

[119] None of the forgoing is in opposition to the point, raised by the ever-helpful Levin, that mathematicians do speak of "computer experiments" in certain contexts. For example, one might have a highly complicated equation, such that one was unable to calculate what value for x yields a particular range of values for y. One might use a computer to calculate the values of y for the inputted values of x. However, such usage presents no problem for the apriorist. Firstly, the computer is used in what even Tymoczko agrees is a harmless way—mere calculation. It is used only as a convenient way of counteracting human error, and does not provide any significant new results. As such, its use is entirely eliminable, replaceable by human mathematicians tracing a priori steps. Moreover, as shall be argued below, if the methodology used by the computer were to model a priori reasoning, then the mere fact that a computer was used would give us no reason to say that the proof was anything other than a priori.

One way of unpacking this is to say that it is in the nature of such a proposition and its relationship to mathematical reality, that experience does not enter into the warrant for that proposition.[120] To say that one could know the 4CT a priori is to say that if one had but time enough, one could, using such processes, discover the truth. All of this is consistent with saying that, in fact, the 4CT is not actually known a priori. Still, the 4CT is an a priori truth because it is the kind of question that is solved using a priori processes. From the perspective offered by this account, Tymoczko errs in concluding from the fact that the 4CT is not in fact known a priori, that it is not an a priori truth.

As this account classifies the 4CT as an a priori truth, we are left with no difficult questions about what it is that gets registered by experiments in mathematics, or what the role of such experiments might be. The unsavoury ends facing those who would classify the 4CT as an experiment ought themselves to lend credence to an account of the a priori which avoids these ends. The plausibility of the account can be further appreciated by noting a distinction which it preserves, but which would be obliterated if we were to follow Tymoczko in assigning a posteriori status to computer-assisted proofs . Consider a set of numbers so large that, while a computer could inspect each of its elements, no human mathematician could. Now consider two computer-assisted proofs, both of which show that, for this set of numbers, the number two is the only even prime. One computer proves this by churning through all the even numbers, and dividing each by two, thus showing that each (apart from 2) is divisible by a number other than itself and one. The other computer proves it from the definitions of

---

[120] Katz, in conversation.

what it is to be prime, and what it is to be even, and, from the nature of these concepts, concludes that there cannot be any other even primes. Both computers have provided warrant for the proposition that the number two is the only even prime in the set. The kind of warrant they have provided is, however, quite different. The first computer has gone through a process somewhat akin to empirical observation. It has examined each member of a given set of objects to see which of those objects possess certain features, and concluded that only the number two possesses the relevant features.[121] Moreover, its conclusion holds only for the set of numbers actually examined. The second computer, on the other hand, has provided an a priori proof. It has used the fact that, with the exception of the number two, the concepts 'prime' and 'even' are necessarily mutually exclusive. Its conclusion holds of all numbers, not just of those in the set inputted.

Of course, this simple example is very different from the case of the 4CT. The second computer's proof is surveyable, for example, and the first computer's proof is nothing beyond mere inspection. The point at issue, however is independent of such differences. It is merely that the nature of the warrant behind a claim is determined not by whether or not that warrant is provided by a computer, and not by whether it is surveyable or not. It is determined rather by the type of process(es) standing behind that warrant. A priori methodology, regardless of the physical nature of the entity which runs through that methodology, provides a priori warrant.

[121] Of course, the computer cannot empirically examine numbers, since these are abstract. Instead, the computer examines the symbols assigned by the program to the numbers.

In light of this, consider, Tymoczko's (1979, 58) remark that "we must admit that the current proof is no traditional proof, no a priori deduction of a statement from premisses." It is true, in one sense, that the 4CT is no traditional proof. However, when we bear in mind that an a priori truth is one which can be discovered using a priori methodology, Tymoczko's gloss on this point is seen to be mistaken. An a priori deduction of a statement from premisses is exactly what the computer produces. Unsurveyable by any one mathematician, it is true, but we have already answered this point. The truth proven is arrived at using exactly the kind of steps that a human mathematician would go through if he could live long enough. At no point do any of the premisses make use of any empirical content. The truth of the 4CT is proven, that is, entirely via a priori methodology.

At the beginning of this section, it was claimed that a the proof of the 4CT would only present a genuine threat to mathematical apriorism if (a) the "experiment" were only properly describable as a posteriori, and (b) the proof were not discoverable using purely a priori processes. Neither (a) nor (b) is satisfied by the 4CT. The 4CT is not really analogous to a physical experiment, and there is a way to describe what is going on in the appeal to computer which avoids the problems which are raised when we try to see it as an experiment, and which further underlines the a priori nature of the proof of the 4CT. Moreover, there is a more natural analogy for the 4CT than that between it and an experiment. To see it, consider our conclusions from the foregoing paragraphs. The computer does what human mathematicians would do if they could but live long enough. It executes a program designed to model the kind of reasoning a human mathematician would go through. Pace Tymoczko, it deduces a conclusion from premisses, using a priori processes. None of the steps have any empirical content. The role of the computer in the proof of the 4CT is, in short, not so much analogous to a physical experiment as it is to that of a robot mathematician.[122]

To better fix the nature of the analogy, consider how chess masters use both human coaches and chess-playing computers. Both are used for the same purposes: to sharpen the skills of, and to provide practice for, the chess master. As far as the chess master is concerned, it makes no difference whether he learns from human or computer, once he learns. Moreover, both human and computer "coaches" manage to play this role by virtue of using similar methods; both have a large number of "stored" games, moves and strategies in their respective memories. Similarly, in mathematics, robot (computer) mathematicians and human mathematicians use the same (a priori) methodology, and play the same role (i.e., discovering mathematical truths or theorems). True, one difference between the chess case and the mathematics case is that, so far as the chess master is concerned, it makes no difference whether what he learns is a priori or a posteriori. As these epistemological categories are unimportant to him, he can regard the computer "coach" as no different from the human. Epistemological categories are important when it comes to knowledge of mathematics, however. If the epistemological status of a particular piece of mathematics may be determined by whether a human or a computer does the mathematical work, we cannot be so blasé as the chess master. And Tymoczko's point is that the case of the 4CT shows that whether a computer or human does the mathematical work can make a difference as to whether particular proofs are a priori or a posteriori. However, once the work of undermining Tymoczko's reasons for

[122] I owe this analogy to Jerrold Katz.

this claim is completed (over the course of the remaining sections), this difference between the chess case and the mathematics case disappears. At this point, one can legitimately seek to elucidate the robot/human mathematician analogy in terms of the computer/human chess coach analogy.

### 3.2.3

### Hardware Errors

Tymoczko (1979, 74) claims that there are three further features of computer assisted proofs which make it natural to describe them as being a posteriori . These are, respectively, the possibility of a hardware error, the possibility of a software, or programming error, and the lack of certainty provided by computer proofs. I shall examine each of these in turn.

Suppose we could increase human mathematical ability with a pill which increased neural activity in the parts of the brain associated with mathematical reasoning. Human mathematicians who took such a pill could now move through proofs much faster, to the extent that they could now survey the entire proof of the 4CT. Suppose further, that all mathematicians who took such a pill reported that the proof was in fact sound. In such circumstances, we would not, I think, have any sound basis for saying either that the verdict of these mathematicians was the outcome of an experiment, or that the theorem was not known a priori. There is, moreover, a conception of a priori truth in the literature, due to Tamara Horowitz, which captures the intuition here (and which can be seen as a way of unpacking the conception of a priori mentioned a few pages back—i.e., that it is in the nature of an a priori proposition and its relationship to mathematical reality, that experience does not enter into the warrant for that proposition). According to it, a truth is a priori iff:

> some member of some [epistemological] type can know p independently of experience, and, for all types and all metaphysically possible worlds, if a member of a type in a world can know p independently of experience, then any possible member of that type can likewise know p independently of experience in any world in which it exists.[123]

The pill-enhanced mathematician knows the proof of the 4CT independently of experience, and mathematicians in our world are possible members of the type of such enhanced mathematician. Thus, at least one defensible conception of the a priori accords with the intuition that a pill-enhanced mathematician could know a priori that the proof of the 4CT was sound, and thus with the claim that the 4CT is an a priori truth.

There is one consideration which might make it seem as though we should say that our pill-enhanced mathematician does not know the truth of the 4CT a priori. We might suggest that our belief that the proof was sound rested on an empirical assumption about the reliability of brains which had been affected by the chemicals contained in the pills. But this consideration actually makes the apriorist's point for him. For it is simply to say that when we accept proofs produced by human

[123] Tamara Horowitz, "A Priori Truth," The Journal of Philosophy, Volume 82, Issue 5 (May, 1985): 231 – 232. Horowitz goes on to reject this conception of the a priori. Her rejection of it is based on a counterexample which conflicts with the claim (which she takes to be central to the a priori) (C): A priori propositions express facts such that if those facts did not obtain, nobody would be able to have precisely the concepts ingredient in the propositions. Crucially, Horowitz's counterexample involves a contingent a priori truth. As she herself acknowledges (233), her counterexample shows only that (C) does not hold of contingent a priori truths. One could, therefore, consistently hold a version of (C), restricted to necessary a priori truths, together with the conception of a priori truth outlined in the text.

mathematicians, we assume their brains are functioning appropriately. And if this suffices to show that knowledge of theorems is empirical, one might very well ask just what the entire debate between mathematical empiricists and apriorists was about anyway.

That brains are functioning correctly is a background condition for warrant (be it a priori or a posteriori). It is not itself (as Burge (1998, 6) emphasises) part of the warrant for the belief in question. Put another way, a correctly functioning brain is something which enables a justification (or proof) to be understood. It is not itself a premiss in that proof or justification. Similarly, the proposition that the computer is functioning properly is not itself part of the proof of the 4CT. And, if it is not, the presence of a "By Computer" lemma does not represent an appeal to a physical experiment, or to the internal physical workings of the computer. Since it does not, it presents no problem to mathematical apriorism.

### 3.2.4

### Programming Errors

I shall be brief with Tymoczko's two remaining worries. Firstly, he mentions the possibility of an error (or bug) occurring in the computer program itself (i.e., a software, rather than a hardware error). Given this possibility, Tymoczko (1979, 74) concludes that the 4CT does not have the same degree of reliability guaranteed by traditional proofs. But this is a simple matter for the apriorist to deal with. As Burge (1998, 8) observes, computer programming languages are just special types of mathematical (or logical) languages. The reliance on these by the 4CT does not render it a posteriori, any more than does the reliance of traditional proofs on other formal languages render them a posteriori. There may be errors in the program language, just as there may be errors in the use of formal languages by mathematicians and logicians working in traditional areas. In either case, if an error is serious enough, the "proof" will not yield knowledge. Moreover, in the cases of proofs offered both by computers and by human mathematicians, the method we use to weed out such errors is the same: the proofs are checked, a priori, by other competent reasoners.

### 3.2.5

### The Argument From Uncertainty

There is one final feature appealed to by Tymoczko to make his case that the warrant provided by the computer to the 4CT is a posteriori. This arises from Tymoczko's (1979, 77-8) reflection on the fact that our knowledge must be qualified by the uncertainty of our instruments. There are, he tells us, truths from electrical engineering which have a higher degree of certainty than the 4CT. Tymoczko takes our lack of certainty regarding our instruments to cast a shadow over the claim that the 4CT is an a priori theorem. 'Certainty' here may mean one of two things, but on either reading, Tymoczko's claim is mistaken. On the first reading, 'certainty' means objective certainty. To say, in this sense, that the truths of mathematics are certain is to say that their truth is objectively certain, regardless of what we happen to think about them.[124] It is objectively certain, for example, that seven times seven is forty-nine, and this truth

[124] For more on the notion of objective certainty, see Katz (1998), 63-64.

would be certain even if everyone came to believe otherwise. In this sense of certainty, Tymoczko's claim is false. The truths of mathematics are objectively certain, and thus cannot be less certain than any truths of electrical engineering. On the assumption that the proof of the 4CT is a truth of mathematics, it cannot be less objectively certain than any other.

On the other hand, by 'certainty' Tymoczko may (and probably does) have in mind our degree of confidence in the truth of the 4CT. In this sense, he is correct to say that there are truths of engineering with a higher degree of certainty than the 4CT. But the admission is harmless to apriorism, since our confidence regarding our beliefs is not the issue. If it were, then sentences like 'There have appeared to have been brown dogs' would be a priori. As I have stressed throughout this chapter, whether a proposition is known a priori or not has to do with its warrant, not our confidence in its truth. Moreover, again as has been stressed here, since a priori warrant is defeasible, a priori knowledge is compatible with uncertainty. Neither the fact that we have less than one-hundred per cent confidence in the 4CT, nor the fact that our epistemological warrant is less than certain, pose a threat to the a priori status of the 4CT.

Finally, given that the uncertainty over our instruments will be uncertainty about either hardware or software, the conclusions of the previous two sections serve also to defuse this worry.

At the end of section 3.2.2, we rejected the analogy between scientific experiments and the role played by the computer in the proof of the 4CT. We suggested a more appropriate analogy would be that between the role of the computer and a robot mathematician. We then attempted to elucidate the analogy between robot and human mathematicians in terms of the analogy between computer and human chess coaches. The problem with elucidating matters in this way was that there was a relevant difference between the two analogies. While it made no difference in the chess case whether one learned from a human or from a computer, Tymoczko's claim was that the case of the 4CT showed that whether a computer or a human were to do the mathematical work would make a difference as to whether particular proofs were a priori or a posteriori. It would, at that point, have been to beg the question against Tymoczko, had we left standing our analogy between the role of the computer in the 4CT, and the role played by computers in chess coaching. Since, however, we have now completed the work of undermining the reasons Tymoczko offered in support of this claim of his, we are now free to elucidate matters via this analogy, and the work of section 3.2.2 is thereby completed.

## 3.3

## Conclusion

None of Tymoczko's attempts to elucidate the threat posed by the 4CT to apriorism survive scrutiny. The proof of the 4CT does not differ from traditional proofs in any way that would change its epistemological status. Nor can this point be used to argue that traditional mathematics is also a posteriori. On the understanding of the a priori offered here, what makes a mathematical truth a priori is whether it is knowable using a priori mathematical methodology. Both traditional proofs, and computer-assisted proofs are a priori in this sense. Of course, if one defines an a priori truth as one the warrant of which is a priori, the question of the nature of that warrant arises. But we have already seen that a priori warrant is fallible warrant which derives from processes which are both reliable and repeatable, the warranting power of which does not depend on any particular character of experience (although the concepts required may originate from the senses). Following a proof is a process of this type, as are other non-basic mathematical processes. The reliabilist thus counts them as a priori.

Similarly, none of Kitcher's arguments against apriorism succeeded. Since, as I remarked at the outset of this chapter, mathematical apriorism is the default position, the absence of any real threat to that position leaves us free to hold that non-basic mathematical truths are knowable a priori.

An account of a priori warrant for basic mathematical beliefs will be given in Chapter Five. The warranting process here will likewise be shown to be reliable, repeatable, and not derived from any particular characteristic of experience. This will complete the account of a priori mathematical warrant. The more immediate task, to be undertaken in Chapter Four, is to answer the charge that reliabilism falls foul of the Benacerraf problem, and thus cannot possibly give a satisfactory account of the non-inferential mathematical process of intuition.

# Chapter Four

## Casullo on Why Reliabilism Can't Solve the Benacerraf Problem

### Introduction

Paul Benacerraf, in his 1973 paper, "Mathematical Truth," presents the following incompatibility. If Platonism is true (as this dissertation assumes), then mathematical objects such as numbers exist, and do so abstractly. Since they are abstract, we cannot causally interact with them. If the causal theory of knowledge (CTK), which at the time appeared to Benacerraf and others to be the best theory of knowledge, is correct then all non-inferential knowledge requires a causal link between the knower and the object known. Hence the incompatibility: our best theory of knowledge and our best theory of mathematical truth cannot both be true.

It is some time since the CTK was widely regarded as our best theory of knowledge. There is, then, some temptation to regard Benacerraf's problem as having died with the CTK. However, as many authors have observed, the problem is a good deal more robust than the epistemological theory. Thus Field (1989, 25-26) writes:

> The way to understand Benacerraf's challenge, I think, is not as a challenge to our ability to <u>justify</u> our mathematical beliefs, but as a challenge to our ability to <u>explain the reliability</u> of these beliefs ... The idea is that <u>if it appears in principle impossible to explain this</u>, then that tends to <u>undermine</u> the belief in

> mathematical entities, despite whatever reason we might have for believing in them.

And Maddy (1990, 43) remarks that even after we have given up on the CTK:

> [T]here will remain the problem of explaining the undeniable fact of our [mathematical] expert's reliability. In particular, even from a completely naturalized perspective, the Platonist still owes us an explanation of how and why Solovay's beliefs about sets are reliable indicators of the truth about sets.

The problem will remain for any theorist who attempts to combine the empiricist requirement that there be, for non-inferential knowledge, some kind of causal interaction between knower and known with a Platonistic view of mathematical reality. And, for those who opt to reject the empiricist requirement, the challenge is to outline how it is that processes which involve no causal link between knowing subjects and the objects of non-inferential knowledge may nonetheless be reliable.

Reliabilism is the direct descendant of CTK. One of the advantages claimed for reliabilism over CTK was its apparent compatibility with Platonism, a consequence of its definition of knowledge as true belief, which is produced by a reliable process. One striking feature of this simple definition is that, on the face of things, no causal connection between knower and known is required. Once whatever (reliable) process it is that produces mathematical knowledge produces a true mathematical belief, that belief is knowledge—regardless of whether the process is one which involves causal interaction between a mind and a mathematical object. Apparently, then, though reliabilism is a direct descendant of CTK, the family resemblance is not great.

Albert Casullo (1992) has argued that this touted advantage of reliabilism is merely apparent. Upon closer examination, Casullo claims, reliabilism, just as much as the CTK, is seen to require a causal link between numbers and ourselves in accounting for non-inferential mathematical knowledge. Since there can be no such link, reliabilism ultimately falls foul of a version of Benacerraf's problem. The main business of this chapter will be to argue that Casullo is mistaken.

Faced with Casullo's extension of Benacerraf's problem, one could, of course opt to maintain reliabilism by jettisoning one's commitment to Platonism. Given my project, this option is not open to me. There are, however, reasons to think that such a move is a bad one, even if it were available to me. As Benacerraf observed, one who claims that numbers do not really exist must provide an answer to the question "So what are mathematical sentences <u>about</u>?" This, as the many objections to the various forms of fictionalism, nominalism, and conceptualism indicate, is no mean task. The Platonist has an easy answer here. Another reason why I think eschewing Platonism would be a bad way to defend reliabilism is that reliabilism is, to put it mildly, a controversial epistemological theory. Tying it to one side in the dispute over Platonism merely makes it more controversial. If, on the other hand, we can show that reliabilism can be accepted whatever the truth about numbers, it would be on firmer ground.

The reliabilist might instead opt to maintain both his Platonism and his reliabilism by accepting an empirical holism, combined with a Quinean indispensability argument for the existence of mathematical objects. There are a number of problems with this strategy. Firstly, the epistemology described by empirical holism is committed to the revisability of all statements. It thus will not yield necessity,[125] and it has seemed to many philosophers that the truths of mathematics are necessarily true. The Platonism such a strategy offers, then, is rather anaemic by comparison to the traditional variety. Secondly, the approach avoids the Benacerraf problem at the cost of making even the most basic mathematical knowledge highly inferential. Mathematical objects are posited, and beliefs about them justified, via the indispensable role such objects play in physical theory. But it is at least prima facie odd to think that a belief that 2 plus 2 equals 4 is such a highly theoretical belief, and that it depends, for its truth, on theories in the physical sciences, most of which are more abstruse than it. Thirdly, as Quine himself acknowledges, the indispensability argument works only for mathematics that actually is genuinely indispensable to the physical sciences. Mathematics without application is, for Quine, an "uninterpreted system."[126] Finally, as I documented in Chapter Two, there are questions about the coherence of empirical holism itself (Katz 1998, 73; Warner 1989, 46-47). Adopting this approach thus represents an extremely problematic strategy for the reliabilist.

Accordingly, the strategy to be adopted here is to argue that reliabilism, contrary to what Casullo urges, is compatible with the existence of basic belief forming processes which yield warranted, true beliefs about an abstract mathematical reality.

---

[125] One might question the assumption that revisability would imply contingency, since the former is epistemic, whereas the latter is metaphysical. However, given that, for any proposition that we know (or believe) to be necessarily true, revising that proposition (i.e., assigning it a truth value of 'false') would not be rational. So, since holism is committed to the rational revisability of all propositions, it cannot allow that we might know that some propositions are necessarily true. Thanks to Michael Devitt for raising this point.

[126] See his "Success and the Limits of Mathematization," in Quine (1981), 151. Resnik (1997, 146-7) attempts to include unapplied mathematics by means of possible future applications. I argued in section

The outline to my response to Casullo, then, is as follows. Casullo considers several preliminary forms of reliabilism, and finds each of these vulnerable to counterexamples found in the literature. He then sketches an account of what he thinks reliabilism would have to look like in order to avoid such counterexamples. This form of reliabilism, Casullo argues, requires a causal link between our a priori mathematical intuition, and mathematical reality. It therefore falls victim to the Benacerraf problem. I will defend each form of reliabilism considered by Casullo. I will, for each counterexample that he presents, give reason to think that it does not present a real problem for the version of reliabilism which Casullo is considering at that point. Finally, using what Casullo considers the most plausible form of reliabilism, I shall argue that it does not require any causal link between us and mathematical reality, and thus that it avoids Benacerraf's problem. If this strategy is successful, one consequence will be that the weaker forms of reliabilism with which Casullo begins, will in effect have been shown to escape all of the criticisms which he levels at reliabilism.

## 4.1

## Casullo's Arguments

Casullo's discussion focuses on the notion of justification. There are, broadly speaking, two notions of justification in the literature, one internalist, the other externalist. The externalist notion is here represented by process reliabilism, according to which, a <u>belief</u> is justified iff it is produced by a reliable process. For this kind of justification, it is not

---

2.2 that his attempt to extend the indispensability argument will not suffice to warrant all unapplied mathematics.

required that a believer be able to cite, or even have access to, the justification of that belief. On the internalist notion, a believer is justified in holding a belief iff he stands in some suitable first-person accessible relation to the evidence for that belief, or if he may properly be said to be, in some sense, default justified (as might be the case with canonical examples of visual beliefs, for example). For current purposes, the main difference between these two notions of justification is the difference between what it is for a belief to be justified, and what it is for a believer to be justified in holding a particular belief. Accordingly, I shall refer to these different notions as belief-justification, and believer-justification. There is then the further question about how these different notions of justification relate to the notion of knowledge. Believer justification is neither necessary nor (when added to true belief) sufficient for knowledge. That children and certain animals can know things, without being able to offer justification for their beliefs, shows that such justification is not necessary. That the visual beliefs of the scientifically unsophisticated can be justified further serves to show this.[127] The fact that Cartesian demons and brain-vat-scientists might induce in their subjects true beliefs, which these hapless dupes then order in an epistemically responsible fashion, shows that believer-justification is not (when added to true belief) sufficient for knowledge.

One might object at this point that internalist believer-justification is important in its own right, even if it is neither necessary nor sufficient for knowledge. It is important because ordering our beliefs according to whether we have good reason to

[127] Goldman's (1967) example of how one can know the date of Lincoln's death even if one cannot produce one's justification for this belief further illustrates this point.

hold them is the only way we can legitimately hold that our beliefs are true. This is true, but harmless. Reliabilism, as I argued in Chapter One (section 1.3.3), can account for why it is desirable to believe all and only those things for which we have evidence: because believing in accordance with evidence is a reliable process.

Following Plantinga (1993, 3), I shall use the term 'warrant' as a placeholder for whatever it is which must be added to true belief to yield knowledge. Warrant, then, is not believer-justification. On the other hand, it might be belief-justification (though of course, it might not be).

With these terminological considerations in hand, we can examine critically Casullo's arguments. The first form of reliabilism that Casullo discusses is due to Penelope Maddy:

(M) S's belief that p is (noninferentially) justified if it results immediately from a reliable process (Maddy, quoted in Casullo, 1992, 571)

(M) is, as formulated, an account of belief justification, not of believer-justification. [128] Casullo (1992, 572) offers three distinct difficulties for (M). Where R is a reliable process, and p is a true belief produced in S by R, (M) will be inadequate, he claims, when any of the following conditions obtains:

1. S has good reason from another source to believe that p is false;

[128] I say "as formulated," because it is not important, for present purposes, whether this is how Maddy actually intends (M) to read.

2. S bases her belief that p on the belief that p is produced by R but also has good reason to believe that R produces a preponderance of false beliefs in her case;
3. S bases her belief that p on the belief that p is produced by R but also has good reason to believe that there could not exist a cognitive process such as R.

Let us take each of these worries in turn. Suppose that S believes that p, where p is some basic mathematical truth. Suppose further, that S's belief that p is produced by her reliable mathematical intuition, even though S is not in a position to offer reasons for her belief, and in fact has reason to think that p is false. In spite of this, S continues to believe that p. What is the epistemic status of S's belief that p? As the reliabilist sees things, since S's warrant for p cannot possibly be undermined, the process which led her to believe that p justifies the belief that p. That is, S's belief that p is justified, although she is not believer-justified in holding it. Casullo's case cannot show S is not belief-justified, so, in effect, it depends on assuming that believer-justification is something which the reliabilist account needs to capture. The reliabilist, on the other hand, denies that an adequate theory of knowledge must explicate the internalist notion of believer-justification. For reasons already rehearsed, believer-justification is neither necessary nor (when added to true belief) sufficient for knowledge.[129] (Though the reliabilist might offer an account of believer-justification which explains the desirability of this

[129] One might at this point ask why it would be in any sense warranted for anyone to believe some p, in the absence of believer-justification. Simply because, built into the nature of cases like Casullo's, and Bonjour's, is the condition that there are processes more reliable (in the given situations) than believing in accordance with available evidence. In fact, by the reliabilist's lights, it is the very fact that internalist believer-justification is not a necessary condition on knowledge that allows the clairvoyants in such cases to have knowledge, when they lack this type of justification.

notion in terms of its reliability, he will see this as a different project. The distinction is that between a theory of believer-justified belief, and a theory of knowledge).

Moreover, in the context of Casullo's article, the fact that reliabilism does not account for internalist believer-justification is irrelevant. Casullo's article is an attempt to extend Benacerraf's puzzle about mathematical knowledge to reliabilism. Why should a demonstration that reliabilism cannot account for a conception of justification which is explicitly rejected by reliabilism be thought to be relevant to whether reliabilism can account for mathematical knowledge?[130] Why hold that because a theory does not accord well with internalist justification, it cannot offer an account of mathematical knowledge? Is it not an option that a theory might offer an account of warrant, rather than internalist justification, and thereby answer Benacerraf?

Even if we focus on believer-justification, it is not at all clear that Casullo's case is a genuine counterexample. Where p is some mathematical truth, and S continues to believe that p, in spite of her having good reason to believe not-p, two possibilities arise. Either she has considered her good reason to think not-p, and found it wanting, or she has ignored it. The first of these cases is straightforward. S has examined two conflicting sources of (prima facie) believer-justification with respect to p, and, happily, has opted for the right source. She therefore rejects the "good" reason for believing not-p, and continues in her believer-justified belief that p. Alternatively, S continues to believe that p because she simply ignores the "good" reason to believe that not-p. Remember, the p in question is supposed to be true. Moreover, as it is a mathematical truth, it is necessarily true. Subsequently, there can be no possible or actual defeaters for

[130] See also the responses given to Bonjour cases in section 1.3.1.

this belief. Hence, refusing to consider an alleged defeater for p cannot undermine her believer-justification for p. That one can be believer-justified in one's true mathematical beliefs even if one does not consider what would, if one did consider them, appear to be defeaters, seems independently plausible. We would want to say that most people are believer-justified in their belief that 1+1=2, even though they might not have encountered, and so would not be in a position to refute, the bogus, but initially compelling "proof" that this is not true.

Moving on to Casullo's objection (2), two questions arise. How good is S's "good" reason for believing that R produces a preponderance of false beliefs in her case? Given this undermining evidence, does she still believe that p? Take the first of these. As with (1), it may be that S sees her warrant for p as being stronger than the warrant arising from her "good" reason to believe that R is unreliable. She may then, while still being epistemically responsible, discard this "good" reason. She may, on the other hand, decide (mistakenly) that R is too unreliable, and discard her p-belief. Neither of these options are problematic. On the other hand, S may retain her belief in p, while acknowledging (incorrectly, as it turns out) that she has good reason to reject the deliverances of R. What follows? Not that S does not know that p, nor that S is not in any sense justified in her belief that p. What follows is only that S is not believer-justified in holding her belief that p. Since p is true, and produced by the reliable process R, S is still, by reliabilist lights, belief-justified in her belief that p. Her belief, again by reliabilist lights, is thus warranted, and counts as knowledge. Without an argument that internalist justification is the only kind of justification, or that it is necessary for knowledge, nothing follows about whether S's belief is not in any sense justified, or whether S knows that p. The reliabilist who adheres to (M) can claim that although S is not believer-justified in her belief that p, she nonetheless knows that p because she is belief-justified.

The challenge posed by (3) can be answered in either of two ways. Firstly, one could argue along similar lines as were employed in answering (1) and (2), that (3) shows only that there can be cases where one's belief is justified, even though one is not believer-justified in holding the belief. Alternatively, one might offer the following case. Consider Hartry, mathematician by day, philosophical nominalist by night. Qua mathematician, Hartry forms many true mathematical beliefs. Some of these are formed via proofs, but some—the more basic ones—are formed via his ultra-reliable faculty of intuition. Qua philosophical nominalist, however, Hartry reads, and indeed develops, several weighty arguments against the proposition that there exists a faculty of intuition. So weighty does Hartry consider these arguments, that he rejects this proposition. Now, in coming to believe that there is no faculty of intuition, does Hartry thereby cease to have believer-justification for his mathematical beliefs? To claim that his philosophical scruples undermine his mathematical believer-justification flies in the face of the evidence. He can find his way through complicated proofs, spot holes in the proofs of others, tell sound reasoning from fallacious reasoning, and argue with the best of them. We can even suppose he goes on to discover an extremely important new axiom, perhaps revolutionising the whole of mathematics. What sense remains of the claim that none of his beliefs are believer-justified, simply because he does not believe in intuition? (3), then, at least when R is taken to be a reliable process of mathematical intuition, is simply not a problem for (M).

In fact, a Hartry-style case would also defuse (2). Consider a version of Hartry who believes that he has intuition, but that his intuition is unreliable. Perhaps he has been convinced by Wittgenstein's point that if intuition can lead him right, it can also lead him wrong. As it turns out, Hartry's faculty of intuition, contrary to what Hartry thinks, is actually extremely reliable. As before, we can make Hartry as talented as we like. Are we really to say that he lacks believer-justification for his mathematical beliefs, simply because he incorrectly assesses the reliability of his intuition? The point of both Hartry cases is that having believer-justification for one's mathematical beliefs—and indeed having mathematical knowledge—is independent of one's philosophical beliefs about the ontology and epistemology of mathematics. If this point is correct, then one does not lack believer-justification, or knowledge, simply because one lacks (or has mistaken) higher-order beliefs about the faculty which produces one's mathematical beliefs.

One can then, hold fast to (M), provided one is willing to hold (i) and (ii).

(i) Believer justification is not the only kind of justification there is. One's beliefs can be belief-justified without being believer-justified.

(ii) One can have knowledge without having believer-justification.

One could also hold (M), and argue that it is consistent with requiring believer-justification if one were also to hold (iii) and (iv).

(iii) One can be believer-justified in believing necessary mathematical truths, even if one does not consider what would, if one did consider them, appear to be defeaters.

(iv) One can have knowledge of mathematical truths, (and even be believer-justified in holding those beliefs) even if they are produced by a faculty in which one does not believe.

Casullo (1992, 572), taking each of (1)-(3) to constitute a refutation of (M), considers another version of reliabilism, this time due to Goldman:

(G) S's believing p at t is justified iff:

(a) S's believing p at t is permitted by a right system of J-rules, and

(b) this permission is not undermined by S's cognitive state at t[131]

where a right system of J-rules is one which permits basic reliable psychological processes. Clause (b) serves to rule out cases where the permission granted by (a) is undermined by other beliefs of S's (as in cases (1)-(3) above). The language of (G) suggests that Goldman is explicating the notion of believer-justification, rather than belief-justification. As we shall see, however, whether this is so depends on how clause (b) is to be understood.

[131] Alvin Goldman, Epistemology and Cognition (Cambridge, Ma.: Harvard University Press, 1986), 63.

Supposing, for the sake of argument, that Casullo is correct in rejecting Maddy's formulation of reliabilism, what can be said about Goldman's? Casullo (1992, 572) invokes Bonjour's case of Norman against (G). Norman, recall, is a reliable clairvoyant, possessing no evidence for or against the proposition that he is clairvoyant. His clairvoyance regularly causes him to form true beliefs about the whereabouts of the President. According to Bonjour, Norman satisfies both (a) and (b), and thus, since Norman is not justified in his President-belief (according to Bonjour and Casullo), G is refuted. Goldman's (1986, 112) response is to argue that the case does not really satisfy (b). He argues that Norman ought to reason that if he genuinely were psychic, he would surely find some evidence for this. Since he has no such evidence, he ought to conclude that he has no psychic powers. Norman is thus ex ante justified in believing that he is not psychic, and this undermines his psychic President-belief. This latter belief, thus does not satisfy condition (b).

Casullo (1992, 575) uses this response to divide reliabilism into weak and strong forms, according to whether it requires that a cognizer have beliefs about the process which generates a belief in p in order to be justified in believing that p.

(WR): In the absence of undermining evidence, the mere fact that R is a reliable process and produces in S a belief that p is sufficient to justify S in believing that p.

(SR): In order for S to be justified in believing that p, where p is produced by some process R, S must be justified in believing that p is produced by R.

(SR) is clearly an explication of believer-justification. (WR) is somewhat ambiguous. Since it places no conditions on justification beyond mere reliability, it seems to be explicating belief-justification. On the other hand, (WR) states that reliability is sufficient to "justify S in believing that p," and thus it seems to be an account of believer-justification. If (WR) is intended as an explication of believer-justification, it is inadequate, since such justification, by definition, requires more than mere reliability. Thus, I intend to read it as an explication of belief-justification.

Casullo (1992, 575) claims that answering Bonjour requires abandoning (WR) in favour of (SR). In effect, this is to claim that a condition of adequacy on a reliabilist response to Bonjour is that it incorporate an explication of believer-justification. Whether this is true, however, depends on how we understand Bonjour's challenge. The Norman case can be understood as a challenge to reliabilism in three distinct ways. Firstly, it can be understood as a counterexample to the reliabilist conditions on knowledge. Understood this way, the case is that, since Norman satisfies all of the conditions reliabilism places on knowledge, but yet does not have knowledge, his case refutes the reliabilist analysis of knowledge. This claim is answered, as I argued in Chapter One (section 1.3.1), simply by noting that, since it assumes that access to the justification of one's belief is necessary for knowledge, it begs the question against the reliabilist. So, if the challenge is understood in this way, then Casullo is mistaken to claim that answering Bonjour requires abandoning (WR).

The two other ways in which the challenge can be understood both have to do with how reliabilism handles justification. We have distinguished two notions of justification in the literature, one internalist, the other externalist. The challenge posed by the Bonjour cases differs depending on which way we understand justification. Bonjour cases might be interpreted as cases which refute reliabilism by showing that, while Norman satisfies the conditions (WR) places on externalist belief-justification, his belief is not belief-justified. Bonjour's case does not show this, and was not intended to. Norman's beliefs, by definition, satisfy belief-justification.

Finally, consider the Norman case, construed as an allegation that (WR) cannot handle believer-justification. This allegation, so far as I can see is correct. If clause (b) is understood along the lines of (WR), then (G) does not capture believer-justification. Thus, if Casullo's conclusions about (M) were correct, (G) would be likewise inadequate. If (b) is understood along the lines of (SR), then (G) has moved a large part of the way towards becoming an account of believer-justification, which Casullo implicitly takes to be a condition of adequacy. I argued earlier, in the discussion of (M), that believer-justification was not a condition of adequacy on an account of justification. I therefore do not accept Casullo's conclusions concerning (M), and therefore do not consider the failure of (WR) to incorporate an explication of believer-justification as damaging to that theory.[132]

So, no matter how the Bonjour cases are understood, they cannot show that (WR) must be abandoned in favour of (SR). One can, that is, maintain (WR) if one is also willing to hold that one can be both (belief-)justified, and can know some p, even if one has no beliefs whatever about the process which actually produced the belief that p.

[132] See also my discussion of Bonjour in section 1.3.1, where I argued that Bonjour can be answered within an account which incorporates only belief-justification, and not believer-justification.

Thus far we have seen that the reliabilist can provide answers to Casullo's treatment of both (M) and the (WR) reading of (G). However, for the sake of argument, let us assume that Casullo's objections have forced the reliabilist to accept that he must account for believer-justification—specifically, that he must accept the (SR) reading of (G). Casullo next argues that even (SR) is insufficient to escape the Norman case. The (SR) reading of (G) requires that if S is to be justified in believing that p, S must be justified in believing that his belief that p was produced by a reliable process. Since clause (b) of (G) explicates justification solely in terms of a cognizer's cognitive state, one could satisfy (b), and thus be justified in holding one's belief, simply by refusing to listen to evidence which undermines one's belief in the reliability of the process R. Granted this, Goldman's clause (b) needs to be augmented. Casullo's suggestion is (b*):

> (b*) this permission is not undermined by S's cognitive state at t or other evidence available within S's epistemic community at t.[133]

But there is good reason to think that (b*) cannot be a necessary condition even on believer-justification.[134] Consider once more, the case of Blinky, from Chapter One

[133] (G) now requires of a (believer-)justified belief, p, that S's believing that p be permitted by a right system of j-rules, and that this permission should not be undermined either by S's cognitive state, or by other evidence available within S's epistemic community.

[134] It is, by definition, not a necessary condition on belief-justification, and would certainly be controversial if put forward as a necessary condition on knowledge. I, of course, would reject it as a necessary condition on knowledge, primarily on the grounds that it begs the question against the reliabilist who holds that belief-justification coincides with warrant, and thus, when added to true belief, is sufficient for knowledge.

(section 1.3.1). Blinky, alone amongst the population of his world, can see. Let us suppose that he even has 20/20 vision. There is widespread scepticism in Blinky's community as to the reality of this mysterious "fifth sense." Apart from a few maverick members, the scientific community is unanimous in its rejection of the existence of such a faculty. Nevertheless, Blinky can see. He's not sure quite how the process works, and he has no idea why he alone should have been singled out for such a gift, but Blinky can see perfectly well. Blinky, in optimal conditions, gets a good look at a medium sized object (say, a dog), and forms the belief that there is a dog in front of him. Appropriate to the object Blinky has seen, I shall, at this point, be dogmatic. Blinky is as believer-justified in his visual beliefs as we are in ours. Regardless of what his community thinks of the possibility or otherwise of the existence of a faculty of vision, he is believer-justified in believing there is a dog in front of him.

One might object that the standards of believer-justification are set by, and relative to, the community of which the subject is a part. Blinky might well be believer-justified in his visual beliefs if he happened to live in our community, since we recognize vision as a source of believer-justification. However, the community in which he lives does not recognize vision as a source of believer-justification, and he is therefore not believer-justified in holding his dog-belief. This objection assumes a contextualist notion of believer-justification. It thus assumes a controversial element from a theory of justification which itself is not obviously true. Because it does this, it cannot show that (b*) is a condition of adequacy on believer-justification. An acceptable argument to the effect that (b*) is a necessary condition on a theory of believer-justification must be, broadly speaking, theory-neutral. However, we can still answer the objection by altering the Blinky case so that Blinky is the last human left alive on our planet, and it so happens that he knows nothing of optic science. It seems to me that this Blinky is exactly as believer-justified in his visual beliefs as any of us are in ours. And, since he has no community, the reason for his being so justified cannot be that his community sanctions visual beliefs. If this verdict is correct, (b*) is not a necessary condition on believer-justification. At the very least, unless independent argument is presented to the effect that Blinky is not believer-justified in believing as he does, I take his case to be prima facie evidence that (b*) stands in need of revision.

Drawing the parallel between Blinky and the case of mathematical intuition, it seems to me that one would be believer-justified in one's mathematical beliefs, given that they were formed by a reliable process of intuition, even if empiricists convinced everyone that there was no such thing as intuition. This conclusion can be supported by building on the case of the person who has not considered the misleading "proof" of the falsity of '1+1=2'. We could suppose that everyone else in his community became convinced by this "proof," but, for whatever reason, that "proof" does not come to the attention of this one person. It seems false to hold that such a person would not be justified in believing that 1+1=2. Casullo's (b*) would yield this verdict.

At the very least, we can say that (b*) is extremely controversial as a necessary condition on believer-justification. It seems most convincing when one considers clairvoyance (but perhaps because of our beliefs about the unreliability of clairvoyance in our world). It is not at all convincing for the Blinky cases, and seems, to me at least, unconvincing as a necessary condition on believer-justification in mathematics. What we can say is that Casullo's case that (b*) is a necessary condition on believer-justification not been made.

In any event, Casullo uses (b*) to shape his next counterexamples to reliabilism. In keeping with the strategy of this chapter, let us, for the sake of argument, grant that (b*) is a necessary condition on believer-justification. According to Casullo, two further questions arise.

C1: Can we have justified beliefs that some of our mathematical beliefs are produced by intuition?

C2: Will the justification of this belief survive undermining evidence from the community?

Again, we are here concerned only with believer-justification, not with belief-justification, warrant, or knowledge. Both (C1) and (C2) would be answered positively if we were concerned with (M) or with the (WR) version of (G). The relevant question in considering (C1) and (C2) is whether we could be justified in holding the relevant beliefs. From here on, 'justification' will be understood to mean 'believer-justification'.

In order to argue that (C1) must be answered negatively, Casullo gives us the case of Norma, a person who reliably forms true beliefs about elementary mathematical propositions, via a process of intuition. Norma is naive in the following sense. She has never considered the issue of the justification of her mathematical beliefs, or the question of whether she has a faculty of intuition. She possesses no evidence for or against the proposition that such a faculty exists. Casullo (1992, 579)concludes that

"Norma's epistemic position with respect to her ability to form mathematical beliefs via intuition parallels Norman's epistemic position with respect to his ability to form beliefs via clairvoyance." Casullo (1992, 580) further concludes that Norma's situation is typical of most ordinary folk with respect to their mathematical beliefs, and finally, that the unreflective intuitionist, "i.e., one who forms mathematical beliefs by the process of intuition but who has never critically reflected on the question of which process generates those beliefs, does not have justified mathematical beliefs."

Several responses to this case are in order. The first of these is that since the analogy drawn between clairvoyance and mathematical intuition is a dubious one, the case tells us nothing about justification for mathematical beliefs. Intuition, on the assumption that it exists, is less problematic in at least two respects than is clairvoyance. Firstly, if it exists, and is responsible for our mathematical beliefs, then everyone has it, and is fairly familiar with its operations. This is not so for clairvoyance. This faculty, on the assumption that it exists, is assumed to belong to only a very few people, and thus its operations are quite alien to most of us.

A second problem for the analogy arises out of the respective subject matters in which the controversial faculties trade. Both faculties are alleged to be methods of justification which do not require causal contact between the believer and the subjects of those beliefs. Such contact is not necessary for intuition, since it produces beliefs in propositions which necessarily have the truth values they have. As things could not possibly be other than they are in mathematical reality, sense-experience is not required to justify our beliefs about such a realm; reason suffices. Clairvoyance, on the other hand, is alleged to yield justified beliefs in contingent empirical truths without any causal contact between the believer and the subjects of those beliefs. Empiricists and rationalists agree that this is not possible (at least not in the case of non-inferential beliefs).

Before drawing our moral from this discussion, an aside is required about the appeal to necessity in the previous paragraph. The fact that mathematical truths are necessary serves to allow for justified beliefs about (and knowledge of) mathematics in the absence of a causal connection between us and mathematical reality. It does not, however, render otiose the nature of the process which yields belief in those truths. The reliability of the process of mathematical intuition is not settled by gesturing towards the necessity of the truths in which it produces belief. This is seen by considering a process which produces false beliefs ninety per cent of the time, but produces belief in necessary truths the remaining ten per cent of the time. Such a process is not shown to be reliable by appealing to the necessity of the truth of the true beliefs it produces. If intuition is to yield justified belief in mathematical truths, this will be because of the reliability of its operations over a wide range of truths, and this latter will be dependent on the nature of the process itself, not on the necessity of those truths.[135] The point, then, of the contrast between clairvoyance and intuition, is that no process, including clairvoyance, which produces belief in contingent, empirical truths, could be reliable without causal interaction between the believer and the objects of knowledge. Processes which produce beliefs in a priori necessary truths, on the other hand, provided they are suitably reliable, can yield knowledge in the absence of such causal interaction. My

[135] The nature of non-basic mathematical processes and their reliability has already been discussed in Chapter Three. The nature and reliability of basic mathematical processes will be the subject of Chapter Five.

contention is that intuition, being an appropriately reliable process trafficking in necessary truths, cannot be understood by analogy with clairvoyance.

A further problem for the analogy is that clairvoyance appears to be an incoherent concept, whereas intuition is not. That intuition is not incoherent will, in effect, be shown by the account of it to be given in Chapter Five. That clairvoyance seems to be incoherent is suggested by the following considerations.[136] Since the beliefs formed by clairvoyance are not of a kind which concern the immediate perceptual scene of the subject, but which yet are supposed to be non-inferential, veridical, concerning the external world, contingent and empirical, the only way to understand clairvoyance is by analogy with memory. Memory is the only belief-forming process that produces beliefs which are of this sort. But the analogy between memory and clairvoyance breaks down when we note that the explanation for the veridicality of memory (insofar as memory is veridical) lies in the reliability of the causal connection between memory and previous perceptions (themselves a product of reliable causal connections). Absent a coherent account of backwards causation (together with a reason to think that this is operative in the case of clairvoyance), no such causal link is available when it comes to clairvoyance. There seems, then, no coherent way to understand veridical clairvoyance.

The differences between intuition and clairvoyance noted above show the analogy between true beliefs produced by these processes to be highly dubious.[137]

---

[136] The following considerations grew out of a discussion with Katz.

[137] For the reader who is convinced by the Hartry case above, but still thinks the Norman case is good, the analogy also suffers from the following deficit. Suppose we as a community don't believe in the existence of intuition. Suppose also that it really exists. For such a reader, this would not be enough to render unjustified our belief in 2+2=4. On the other hand, such a reader would consider Norman's clairvoyant beliefs to be unjustified—precisely because we as a community don't believe in the existence

Granted the analogy, however, Casullo (1992, 581) concludes, from his Norma case, that a necessary condition on the justification of beliefs formed through intuition is that one be able to form the following justified belief:

> (I) If S has the faculty of intuition then S forms beliefs of kind K under circumstances C. [138]

He goes on to argue that no-one is in a position to justifiably believe (I). There are two kinds of divergence in the opinions of those who argue for a faculty of intuition, which, Casullo holds, rule out justified belief in (I). Proponents of intuition, we are told, disagree over two very basic points, namely what sorts of beliefs are produced by it, and under which conditions are those beliefs produced. In the absence of some kind of resolution of these disputes, belief in (I) cannot be justified.

Let us begin with the first problem. To exemplify his point, Casullo (1992, 581) notes two divergent accounts of intuition. Mark Steiner's account builds on Ramanujan's "ability to conjecture the most complicated formulas in the theory of elliptic functions, many of which were later proved."[139] Gödel has it that it is the

---

of clairvoyance. This, of course, would be, in effect to deny that (b*) is a necessary condition on justification, which we are here granting for the sake of argument. One could argue, however, that (b*) is less certain than the fact that one does not need to have beliefs about the faculty which produces mathematical beliefs in order for those beliefs to be justified. Thus, in the event of a clash between (b*) and this fact, it is (b*) that should be given up.

[138] Even apart from the fact that Casullo's requirement that one must warrantedly believe (I) if one is to have warranted intuitions is not supported by the Norma case, there are good reasons to reject this requirement. Section 5.1, for example, will offer reasons for belief in the existence of intuition, and the belief that it is responsible for our warranted mathematical beliefs. If the considerations offered there are compelling, one could have reason to believe that one's warranted mathematical beliefs were produced by intuition, without necessarily having a warranted belief in (I).

[139] Mark Steiner, <u>Mathematical Knowledge</u> (Ithaca, NY: Cornell University Press, 1975), 134-135.

axioms that are known through intuition.[140] And, whilst Casullo does not offer it as an example, one might add to the list of possible accounts of intuition, that it is a faculty which grasps basic, though non-axiomatic, mathematical truths. Casullo (1992, 581) notes that while there is no incompatibility in these different accounts of intuition, they "point to radically different places to look for beliefs generated by intuition."

To begin answering this point, we should distinguish between insight and intuition. Insight, whatever else it may involve, is an ability to home in on one's thoughts, and thus to see them in a new light. It is, or involves, an ability to see connections between thoughts, connections which may not have occurred to others. Insight is not limited to mathematical or even logical thought, nor is it limited to non-inferential or basic beliefs in a given area. Ramanujan's spectacular ability, it seems to me, is best thought of as insight, rather than intuition. If this verdict is correct, it is not particularly likely to shed any light on the workings of intuition. If, on the other hand, I am mistaken, and what underlies Ramanujan's ability is indeed intuition, then it is clearly a very highly developed intuition, one able to intuit truths that most others cannot. That being so, it would seem unwise to generalise one's account of intuition from such an atypical example. Either way, beginning with Ramanujan would not appear to be the most profitable strategy for the proponent of intuition.

Gödel's account of intuition is more useful. The axioms of set theory certainly do seem to "force themselves on us as being true." But if one were to take a narrow view of this, and hold both that intuition is a requirement for our having any justified mathematical beliefs, and that intuition is supposed to produce (justified) belief only in

[140] Kurt Gödel, "What Is Cantor's Continuum Problem?" in Benacerraf and Putnam (1983), 484.

axioms, then one is forced to hold that no-one's mathematical beliefs were justified prior to the axiomatization of mathematics.[141] This hardly seems credible.

The third option, not listed by Casullo, is found in Katz (1998): intuition produces belief in non-axiomatic basic beliefs; for example, that four is composite. As formulated, this is in opposition to the narrow reading of Gödel's account given above. However what this third account suggests is that intuition is at least the ability to understand basic mathematical (or logical) propositions and inferences. And once we give an account of intuition in terms of understanding, a path to reconciliation becomes clear. If intuition is to be explained in terms of an ability to understand basic mathematical propositions, then it can produce belief both in the axioms, and in non-axiomatic basic propositions, since both of these sorts of propositions are suitably basic. Indeed, if it is insisted upon that Ramanujan's discoveries were due to intuition rather than to insight, then this account can cover that case too. Ramanujan's highly developed faculty of intuition enabled him to understand connections between propositions without breaking them down into simpler steps. Inferential steps which, for most people, would had to have been broken down into many simpler steps in order for their intuition to comprehend them, were, for Ramanujan's (ex hypothesi) highly developed intuition, readily understandable.

Furthermore, once we look upon intuition as the ability to understand basic propositions, and basic inferential steps, there does not seem nearly so much disagreement amongst the various accounts of intuition on offer. Bonjour's account of a

---

[141] I am not suggesting that Gödel himself held this.

priori justification, for example, has it that "careful and reflective consideration" of a proposition leads one to "see or grasp" the necessity of that proposition. Plantinga suggests that a priori warrant involves finding oneself convinced that a proposition is not only true, but could not have been false. Sosa develops an account based on intellectual seemings, where these are inclinations to believe on the basis of understanding. Hale and Wright offer a Fregean account of mathematical apriorism, where reflection on concepts establishes, a priori, the existence of mathematical objects. Katz (1998, 25-62) has it that reason is sufficient to establish the truth of mathematical propositions. And Peacocke develops his account in terms of possession conditions of concepts. Perhaps Gödel's account might be included here, if the reason for the axioms forcing themselves upon us has to do with our understanding that they must be true.[142]

That there are differences between these accounts (e.g., some reject the contingent a priori, some involve visual metaphors, and there are differences in the amount of knowledge each account classifies as a priori) should not obscure the fact that all of these accounts explain a priori knowledge in terms of understanding. Intuition, then, for each account, will be cashed out in terms of understanding, where that understanding is independent of any particular experiential content. How 'understanding' is understood varies from account to account, to be sure. But psychological theories of understanding will have their differences too, and we would

[142] Lawrence Bonjour, "Toward a Moderate Realism," Philosophical Topics, 23 (1995): 53; Alvin Plantinga, Warrant and Proper Function (New York, Oxford University Press,1993), 105; Ernest Sosa "Rational Intuition: Bealer On its Nature and Epistemic Status," Philosophical Studies 81 (1996): 154; Bob Hale and Crispin Wright (inter alia), "A Reductio Ad Surdum? Field on the Contingency of Mathematical Objects," Mind 103:410 (April 1994): 169-183; Christopher Peacocke (inter alia), "Explaining the A Priori: The Program of Moderate Rationalism," in Boghossian and Peacocke (2000), 255-85.

not, on that basis, claim that they were not attempting to describe something real. Why should there be a double standard for an explication of intuition in terms of the understanding?

What of Casullo's second objection to anyone justifiably coming to believe (I), that of the different accounts of the conditions under which intuition produces mathematical beliefs? Casullo lists Kant's constructivist account, Steiner's (1975, Ch.4) account of intuition as a process of abstraction from perceived or imagined objects, Maddy's view of intuitive mathematical beliefs as pre-linguistic beliefs accompanying the acquisition of mathematical concepts,[143] and Pollock's account of it as being a "seeing that" of relations among abstract entities.[144] He concludes that until some consensus on this issue is forthcoming, no-one can justifiably believe (I).

There is disagreement here for sure, but in the context of my project, it is not harmful. My project, remember, is to reconcile reliabilism with a priori mathematical knowledge, <u>on the assumption that Platonism is the correct ontology of mathematics</u>. On this assumption, then, mathematical intuition is the understanding of basic propositions and inferential steps having to do with a realm of abstract objects, and any account which differs with regards to ontology is (on this assumption) mistaken. Thus, there simply is no disagreement on this issue which matters to my project. One could, of course, pursue a Casullo-type line, and ask how one could come justifiably to believe that Platonism is the correct philosophy of mathematics (and so, come justifiably to believe (I)), but here the general outline of an answer is straightforward, even if the

[143] Penelope Maddy, "Perception and Mathematical Intuition," <u>Philosophical Review</u> 89 (1980): 184-186.

[144] John Pollock, <u>Knowledge and Justification</u> (Princeton, 1974), 318-321.

details are not. One could, justifiably, come to believe that a Platonist ontology was the one which did best justice to mathematical practice and mathematical truth.

Moreover, the fact that there is disagreement about how exactly a mental capacity works does not, by itself, dictate that we cannot justifiably believe in it. If it did, then the proliferation of theories of consciousness ought to imply that we should not believe in that. And, while this conclusion might be welcomed by those of an eliminativist stripe in the philosophy of mind, one could mount a similar argument against belief in vision. One might, that is, argue that one should not believe in vision, since Berkeley-type idealists, phenomenologists, and realists all hold different theories of vision. This argument, surely, no-one would welcome. Casullo, in fairness to him, only sees such disagreement as precluding justified belief where the very existence of that capacity is controversial. However, the first of the Blinky cases lends support to the idea that one can be justified in believing that one has a mental capacity the existence of which is controversial. The fact that the existence of the process which produces a belief is controversial is not sufficient to undermine the justification conferred by that process. Neither is the fact that there are different conceptions of how a given process works. Perhaps the two conjointly might be sufficient to merit scepticism about a putative process, but this case would have to be argued, and Casullo does not argue it. I contend, then, that Casullo gives us no reason to think that one could not justifiably believe (I), once intuition is understood as that mental capacity by which we understand basic propositions and inferences which concern an abstract mathematical realm.

On the supposition that someone could justifiably believe (I), Casullo goes on to argue that this justification will be undermined by evidence available within our epistemic community, and so question (C2) will be answered negatively. The argument, in outline, runs as follows: for one to justifiably believe some deliverance of intuition, p, one's belief that p must also satisfy (b*), but beliefs produced by intuition cannot satisfy (b*). I have argued against (b*) being a necessary condition on justification, and so would reject this argument on those grounds, but even on the assumption that (b*) is in fact a necessary condition, Casullo's point does not follow. Casullo's argument that beliefs produced by intuition cannot satisfy (b*) makes use of causal considerations of the type raised by Benacerraf. (b*) requires that there must not be available evidence which undermines the justification conferred on one's beliefs by the process of intuition. And, according to Casullo, the only non-controversially reliable basic belief-forming processes are causal. There are only two such: perception and memory, although Casullo is willing to allow that introspection might also be an example. Since all of these processes are causal, and belief forming processes which are not causal include astrology, wishful thinking and dreaming, Casullo (1992, 582) claims inductive support for:

(C3) All basic belief forming processes involve the object of belief as a cause of the belief.

Since, then, intuition is not causal, it is, according to (C3), not reliable, and thus (b*) is not satisfied; thus one cannot justifiably believe the deliverances of intuition.

It hardly needs pointing out that an inductive inference which is based on only three instances (only two of which are non-controversial) is very weak. The proponent of intuition thus need not be overly anxious about responding. However, one simple response that he can offer is that the reliable processes listed by Casullo purport to offer information about contingent, empirical truths through causal contact. Thus, one might concede that the argument shows that causal contact is required in order for basic beliefs about contingent, empirical matters to be justified. But this leaves open the question of basic beliefs about a priori necessary truths. It thus leaves open the question of the reliability of intuition.[145]

Moreover, both sides to the dispute recognise that if intuition does exist, it is unique in its workings. It is this recognition that lies behind the charges of mysticism that empiricists delight in levelling at proponents of intuition. Granted its uniqueness then, the fact that all other basic belief-forming processes require causal contact is irrelevant to the epistemic standing of intuition. We knew already that intuition was unlike other processes. And we knew the respect in which it was unlike them: that it is thought to produce justified beliefs without causal contact between knower and known. To offer an inductive argument against intuition, where the induction is based on the very feature which everyone agrees intuition does not have, cannot undermine the epistemic status of intuition. Of course intuition lacks this feature; that is exactly the point.

[145] Nor does Casullo's list of unreliable processes that purport to offer information about contingent empirical truths in the absence of causal contact offer any inductive support for the claim that intuition is not a reliable basic belief-forming process. Both rationalists and empiricists agree that contingent, empirical (non-inferential) truths cannot be known without causal contact. The question is whether some a priori necessary truths can be known without such contact, and this question is no closer to being settled by noting that processes which are alleged to produce belief in contingent a posteriori truths, but which do not involve causal contact, are unreliable.

Casullo has one final objection, again involving causal considerations. He (1992, 583) quotes Hart on conservation of energy:

> Granted just conservation of energy ... you must not deny that when you learn something about an object, there is a change in you. Granted conservation of energy, such a change can be accounted for only by some sort of transmission of energy from, ultimately, your environment to, at least proximately, your brain. And I do not see how what you learned about that object can be about that object ... unless at least a part of the energy that changed your state came from that object.[146]

On this basis, Casullo (1992, 583) claims empirical support for:

(C4) There cannot exist basic psychological processes which generate beliefs about objects which are causally inert.

But (C4) is merely a statement of one of the main tenets of empiricism. Of course, when you learn something, this has to be accounted for in terms compatible with empirical psychology and the brain sciences. So, of course the information gets stored somehow in the brain, which, again of course, is a change in you. But why this has to be the result of a causal process going from the object known to the knower, is left unexplained. Hart's consternation notwithstanding, this is unlikely to impress a rationalist. There is of course some causal process or other when one comes to believe something. And the energy transferred will often be from the immediate environment, but this is all consistent with reason sufficing to know about an abstract realm of necessary truths.

---

[146] W.H. Hart, "Review of Mark Steiner's Mathematical Knowledge," Journal of Philosophy 74 (1977): 125.

Perhaps what bothers those sympathetic to Hart's line is something like the following. Typically, one's mathematical beliefs are initially formed through causal processes. One hears one's mathematics teacher go through the times-tables, one sees a proof written on the board, one devises some method of making sure one's neighbour isn't stealing one's cattle—all of these are causal processes. Here the rationalist will agree as to the genesis of mathematical beliefs, but he will deny that this process accounts for our justification. In the central cases at least, knowledge of mathematical truths will involve understanding the concepts involved in those truths. And this is precisely what is not obtained through the causal processes listed in this paragraph. What we obtain through such causal processes are our mathematical beliefs, and some of the elementary concepts which go to make up the content of some of these beliefs. The justification for these beliefs, on the other hand, comes through the understanding.

To be sure, one can come to know mathematical truths in an a posteriori way, and thus, through a causal process (though not a causal process connecting knower and known). Believing the testimony of a mathematician, without understanding the proof would be one way of so coming to know. But this is consistent both with the claim that a priori knowledge of abstract mathematical truths is possible in the absence of a causal connection between us and numbers, and with the claims that the ways in which mathematicians come to know mathematical truth are a priori and that they do not causally link the knower with the objects known.

So far we have seen that Casullo's claim that any plausible version of reliabilism will be unable to answer Benacerraf's challenge does not withstand scrutiny. There is still a set of worries, closely related to those raised by Casullo, which seem to motivate the claim that reliabilism cannot account for mathematical knowledge. Reliabilism, as I am understanding that term, requires of knowledge only that it be a true belief, produced by a reliable process. And this, the worry might go, is hopelessly weak. Consider the case of an idiot savant, who can give correct answers to complicated mathematical questions, without being able to offer a justification for her answer. Consider Ramanujan, who, while not by any means an idiot savant, could somehow see mathematical truths in the absence of rigorous proofs. Or consider a case where someone consistently gives the correct answer to mathematical problems only after being hit over the head with a stick, and who is unable to offer any justification for his beliefs. In each of these cases, we have true beliefs generated by reliable processes, but we do not have anything that the rationalist would want to call mathematical knowledge.

While reliabilism as I conceive it has to say that S knows that p iff S's belief is true and reliably produced, it does not have to say that S knows a priori that p. This would follow only if p is produced through some a priori process. So the question to be asked prior to answering the above cases, is whether the process is supposed to be a priori or not. If Ramanujan's beliefs were due to a priori intuition, or to some a priori form of insight, the rationalist reliabilist is committed to saying that Ramanujan knows the reliably produced mathematical truths that he believes, and that he knows them a priori. But this is as it should be. Consider Hardy's (1940, 37) description of Ramanujan as having only a rudimentary understanding of mathematical proof. In spite of this, Ramanujan could somehow see truths that were later rigorously proven—and this ability of his was highly reliable. As Robert Kanigel puts it:

> A few of Ramanujan's results were, Hardy could see, wrong. Some were not as profound as Ramanujan liked to think. Some were independent rediscoveries of what Western mathematicians had found fifty years before ... But many—perhaps a third, Hardy would reckon, perhaps two thirds, later mathematicians would estimate, were breathtakingly new. There were thousands of theorems, corollaries, and examples [in Ramanujan's notebooks]. Maybe three thousand, maybe four. For page after page, they stretched on, rarely watered down by proof or explanation ... all their mathematical truths boiled down to a line or two.[147]

Given the description of Ramanujan's abilities, and given that many of the results presented in the notebooks have subsequently been rigorously proved, it would seem churlish to deny Ramanujan knowledge of these results, merely because he himself did not offer formal proofs for them.

The two other cases are different. The processes involved, whether a priori or a posteriori are not only not proofs, they are not even consciously accessible. The responses, however, are fairly straightforward. If the processes are a posteriori, then the cases boil down to Bonjour cases. My answer to these cases then, would be first to point to the possibility of knowing mathematical truths a posteriori, and then refer the reader to my responses to the Bonjour cases. The subjects' beliefs are knowledge, though they lack believer-justification for these beliefs. The knowledge, however, is a posteriori.

On the other hand, the consciously inaccessible processes leading to the mathematical beliefs might be a priori. That is, the kinds of processes in the brain

---

[147] Robert Kanigel, The Man Who Knew Infinity (Scribners (Macdonald & Co.): London, 1992), 203-4.

which lead to the subject's mathematical beliefs might be the same kinds of processes which are at work in paradigm cases of a priori mathematical knowledge—except for the fact that they are not, for whatever reason, available to conscious introspection. As I hold both that the subject has mathematical knowledge, and that what makes a belief a priori or a posteriori is just the kind of process which leads to the belief, I am committed to the possibility of there being a priori mathematical processes which are inaccessible to the believer's consciousness. The savant, then, along with the unfortunate who is being beaten over the head with a stick, has knowledge. The processes are unusual, to be sure, but that should not count against them. If, in addition, the processes are unconscious versions of paradigmatic a priori mathematical processes, then they not only know, they know a priori.

## 4.2

## Conclusion

I have, in this chapter, argued that Casullo holds two mistaken theses. The first of these is that no unadorned version of reliabilism can give even a prima facie plausible theory of justification. The second is that the Benacerraf problem ultimately arises for any plausible reliabilist theory of justification. I take myself to have shown that there is no reason to think that reliabilism is incompatible with a theory of a priori justification for basic mathematical beliefs.[148] I have not, yet, provided a reliabilist account of a priori justification in mathematics. That is the work of Chapter Five.

[148] Whether or not 'justification' must mean 'believer-justification.'

# Chapter Five

## A Reliabilist Account of Mathematical Intuition

### Introduction

The purpose of this chapter is to outline and defend a reliabilist account of mathematical intuition. As throughout this dissertation, I assume that we have mathematical knowledge. I also assume the truth of Platonism. Epistemological concerns to one side, Platonism offers an extremely strong overall account of mathematics. This point seems to be fairly generally acknowledged (Kitcher, 1983, 58, 60; Benacerraf 1983a, 410-412; Hart, 1996, 60 – 61). Platonism preserves the mind-independent objectivity of mathematics, gives a very (some would say the only) satisfactory account of mathematical truth, and allows the intuitive necessity of mathematical propositions. It neither makes the prima facie odd claim that mathematics is a fiction, nor, since it saves the face-value account of the reference of mathematical statements, does it encounter the nominalist's problem of having to reinterpret these statements. Perhaps alone amongst mathematical ontologies, it is sufficiently rich to provide referents for the infinitely many mathematical terms. Further, Platonism happily accepts the ontological commitments of physical theories which quantify over mathematical objects, and thereby helps explain how the sentences contained in such theories can be true.

Given that the theory has such obvious strengths, one might ask why it is so controversial. It seems to me that the problem which bothers most anti-Platonists is the problem of providing a satisfactory account of our knowledge of a realm of atemporal, aspatial, abstract objects, wholly lacking causal powers.[149] Moreover, this problem, again on the assumption of Platonism, is the problem that arises for reliabilism. Even if I am right that Casullo's attempts to show that reliabilism inherits Benacerraf's problem do not succeed, there remains the positive aspect of the problem—that of actually providing a reliabilist account of our basic mathematical knowledge.

Despite the externalist flavor of this dissertation, I will, in this chapter, concentrate on warranting processes the workings of which are first-person accessible. These processes contrast with more bizarre cases, such as the case of one who is able to produce the correct answer to mathematical problems only when hit over the head with a stick, or that of the idiot savant who can offer no justification for her mathematical beliefs, but who nonetheless, reliably gives the correct answer. Consciously accessible processes are, I take it, at work in the more central cases of mathematical knowledge. In these more central cases, the subject typically begins with premisses which he already

[149] There is also the problem of how reference is possible to abstract objects. The locus classicus for this problem is Benacerraf's "What Numbers Could Not Be" (reprinted in Benacerraf and Putnam, 1983). Although this problem is an important one for the Platonist to answer, I think it is fair to say that it does not present as serious a problem for the Platonist as does the epistemological problem. Firstly, Benacerraf showed merely that numbers could not have determinate reference if construed as sets, since there would be more than one type of set which would have equal claim to "really" being the number. If, as my colleague Russell Marcus has argued in an unpublished paper, we give up on identifying numbers with sets, the problem disappears. Secondly, the argument (as Katz observes in "Skepticism About Numbers and Indeterminacy Arguments," in Benacerraf and His Critics, ed. Adam Morton and Stephen Stich, Oxford: Basil Blackwell, 119-39.) is a form of indeterminacy argument, and can be combated in the same way that indeterminacy arguments generally can be combated. That the argument is a species of indeterminacy argument shows that the problem here does not depend on the abstract nature of mathematical objects, since there are indeterminacy worries about reference to concrete objects. This last point, I think, shows why it is the epistemological worry, precisely because it does depend on the abstract nature of mathematical objects, which is the most serious one for the Platonist.

knows to be true (either through earlier proofs, or because of their axiomatic, or otherwise obvious, status), and he follows a proof, the valid form of which necessarily ensures the truth of the conclusion. Certainly these are the cases for which the internalist about knowledge will reserve the honorific 'knowledge'. For one broadly sympathetic to the brand of reliabilism outlined in earlier chapters, however, what distinguishes the typical knowledge-producing processes in mathematics from the bizarre cases, is simply their accessibility to first-person consciousness. One following a mathematical proof is typically able to offer the steps of that proof as warrant for his or her belief, and that warrant is typically accepted as conclusive evidence that the subject knows the conclusion of the proof. One who invariably gives the correct answer to mathematical questions when hit over the head with a large stick, without being able to explain why he believes what he does, has no such first-person access. But if what underlies such a phenomenon was that the person was subconsciously executing the same processes as the conscious mathematician, then my brand of reliabilism holds that the lack of conscious access to the process is the only difference, and that this is of no relevance to the question of whether the unfortunate knows. <u>This</u> question is settled by the nature of the process which led to the belief. Moreover, if the unconscious processes which lead to knowledge in the bizarre cases are a priori (in the sense that they are the same processes as those which standardly lead to a priori mathematical knowledge, except for the fact that they are unconscious) then, as I claimed in Chapter Four, the mathematical knowledge in question is a priori. Thus, although I shall here concentrate on processes which are accessible to first-person consciousness, this is emphatically not meant to suggest that such access is a necessary condition on knowledge. It is simply that I am narrowing the focus to the more typical kind of case.[150]

What is required from a reliabilist account of intuition? Firstly, some reasons for belief in mathematical intuition must be offered. Secondly, a description of what intuition is, and what it does, must be given. It must then be shown that the process of intuition, so described is both repeatable and reliable. In doing this latter, one thereby shows the process to be a warranting one, by reliabilist lights, and further satisfies the Field-Maddy requirement to explain the reliability of mathematical processes.[151] If, further, the process makes no essential use of the particular character of anything which may properly be called experience,[152] the process will yield a priori warrant. As was

---

[150] One merit of this narrowing of focus is that the account of mathematical knowledge which it yields will be more acceptable to one who is sympathetic to a quasi-reliabilist epistemology, but is uncomfortable with allowing as knowledge cases where the subject has no access to the reliable process which produces his belief. Adding a condition such as Casullo's condition (b*), discussed in Chapter Four (4.1), would limit the class of beliefs counting as knowledge to those reliably produced beliefs which also involved some kind of epistemic responsibility on the part of the believer. In effect, (b*) holds not only that S must have no evidence which, if consulted, would undermine his entitlement to continue in his belief, but also that there must not be, in S's community, any such evidence which is readily available. With respect to the "bang on the head" case, in our world, there is readily available evidence which would undermine the subject's entitlement—the well-confirmed belief that subjects do not tend to give correct answers to mathematical problems they could not otherwise solve when hit over the head. Thus, (b*) would rule such a case not to be knowledge, never mind a priori knowledge. Obviously (b*) would have to be fine-tuned to rule out more sophisticated objections, but adding some such constraint to reliabilism would have the result of ruling out these weird cases, thus making it more attractive to the internalist sympathizer.

[151] See the introduction to Chapter Four for their canonical statements of this requirement.

[152] Here, as in Chapter Three, we do not allow the subjective awareness of one's thought processes to count as experience, at least not in any way that is relevant to the a priori/a posteriori distinction. Nothing prevents one from understanding 'experience' in such a way as to include in its scope any event of which one is conscious (e.g., pain, thought, conscious perception). So understanding it, however, obscures an important epistemological distinction between beliefs whose warrant depends solely on intellectual understanding, and those whose warrant depends on the senses. The nature of experience, as it relates to this distinction, is exhausted by processes which take, as input, sensory information, and (mediately or immediately) yield beliefs which depend on some particular characteristic(s) of that sensory information. Thus, if, as shall be shown, the process of mathematical intuition makes no essential use of such

the case with Chapter Three, our definition of a priori processes will not itself be sharpened, but in saying what intuition is, and how it provides warrant, independently of experience, we thereby provide an account of an a priori process which fits our definition, but which is not itself vague. Finally, since we are here dealing with mathematical warrant which is internally accessible, some account of how one could know that one's mathematical beliefs were warranted must also be provided.

Before moving on to considerations which weigh in favor of the existence of intuition, a final preliminary point should be made. As was argued, contra Kitcher, in Chapter Three, a reliabilist account of a priori processes should not require that such processes invariably yield true beliefs. If intuition were to yield a preponderance of true beliefs, without making essential use of any particular characteristic of experience, this would suffice for it to count, properly, as an a priori process.

## 5.1

## Why Think We Have Intuition?

There are several reasons to think that we have intuition. Kitcher (1983, 59), in the course of attacking the concept of mathematical intuition, considers two such reasons. First, there is the argument from Platonist ontology. A realist ontology can account for the truth of mathematics and the mind-independence of mathematics. Only Platonism provides sufficient referents for the infinitely many terms of mathematical language, and (leaving Benacerraf's indeterminacy worry to one side), only Platonism allows for a

---

processes, it is an a priori process. It has already been argued in Chapter Three that mathematical proofs do not involve experience, in this sense, and so are not a posteriori.

standard (i.e., face value) account of reference for mathematical terms. Since these amount to a powerful argument for Platonism, and since Platonism seems committed to the existence of intuition, in order to account for basic mathematical knowledge, we have an argument for belief in intuition. More formally, the argument goes as follows.

1. We have mathematical knowledge (Ass.)
2. Platonism can account for mathematical truth and mathematical objectivity. Only Platonism provides a standard (face value) account of the reference of mathematical terms. Moreover, only a Platonist ontology contains sufficiently many objects to allow all mathematical terms to refer (by examination of theories).
3. So, Platonism is the best overall account of mathematics (corollary of (2)).
4. Thus Platonism is true (IBE, 3)
5. Platonism, in order to fully account for (1), requires a process of intuition to warrant our basic mathematical beliefs.
6. So, we must have intuition (MP, 4,5)

Kitcher's objection to this reason for belief in intuition is that it makes intuition a theoretical posit, and so, all we know about it is what the theory tells us. And, Kitcher (1983, 60) holds, that is not very much. The Platonist's epistemology, is thus "nebulous" (Kitcher, 1983, 61), and, this nebulousness mitigates against intuition counting as an a priori warranting process.

The other reason for belief in intuition considered by Kitcher is the feeling of obvious correctness that we have when we consider simple mathematical truths. We might, he (1983, 61) writes, suppose that " the presence of a feeling of familiarity with basic principles, a sense of their obvious correctness, signals the fact that our belief in them has been generated by an intuition of mathematical reality." But, Kitcher (1983, 61) contends, this familiarity might instead be a result of the conceptual abilities we acquire when we learn mathematics, or from the "indoctrination" of our mathematical youth. Thus, this reason for belief in intuition is likewise inconclusive. Since, then, neither reason for holding that we have intuition survives scrutiny, Kitcher (1983, 64) concludes that the process of intuition could not provide a priori warrant.

However, in examining these two reasons separately, Kitcher has not given the Platonist a fair run for his money. Combining the two gives us a more powerful reason for believing that we have mathematical intuition. Begin with the theoretical posit from Platonist theory. Kitcher's objection to this was that the theory didn't tell us much about it, and so left us unable to decide whether we in fact have mathematical intuition. But now the Platonist only has to overcome the objection that the process he posits is nebulous. The Platonist can begin to do this by moving to the second consideration—our at-homeness with basic mathematical truth. For sure, this might be accounted for in another way, such as one of the ways mentioned by Kitcher, but unless there is some reason for thinking that the Platonist's preferred account of this familiarity is ruled out, he is quite within his rights to answer the charge of unfamiliarity with this response. This then shifts the dialectic to the question of whose account of the familiarity is to be preferred. And this question, since it involves ontological and epistemological aspects, can only be answered as part of the more general question of which philosophy of mathematics is to be preferred. Here, as I outlined above, the Platonist has a strong case.

Of course, to complete his response to the objection that intuition is nebulous, the Platonist will have to offer a description of the process, which tells us what intuition is, and what it does. This I shall do in the remaining sections of this chapter. The point I wish to emphasize for now is just that, since the Platonist posits intuition from within a powerful theory, and since he can answer the charge of unfamiliarity, there is no quick knock-down argument against the two reasons so far given for belief in intuition.

A third reason for belief in intuition goes as follows. Those who reject coherence theories of warrant[153] are left with the following situation. Although belief in some basic mathematical truths may be due to lucky guesses, or to having simply been told the correct answer, at least some of our basic mathematical beliefs are warranted. Since they are basic beliefs, any warrant they are possessed of must be immediate. Since these basic beliefs are intellectual rather than perceptual, the relevant immediately warranting process must itself be intellectual rather than perceptual. And the only candidate for an immediately warranting intellectual process is intuition. Inferences, if

[153] See Chapters 1 and 3 of Susan Haack's Evidence and Inquiry for an excellent, and to my mind, fully convincing, attack on coherentism. There, Haack deals both with simplistic versions of the theory, and with sophisticated versions such as that offered by Bonjour in The Structure of Empirical Knowledge. The problem faced by all forms of the theory is that since a pure coherentism allows no warranting role for inputs from the world, it cannot offer a satisfactory link between the "warrant" provided by coherence, and the truth of a belief thus "warranted." The only way round this difficulty is to allow some warranting role for an input from the external world, as Bonjour does. However, once such an input from the world is allowed to play a warranting role, the theory of warrant, whatever else it may be, ceases to be a coherence theory. Bonjour himself eventually admits that his theory might not be a coherence theory, but dismisses this as an objection to his theory by saying that he is not interested in the label 'coherence theory', just in a correct theory of warrant. This is, of course, a perfectly acceptable move on Bonjour's part, but it in no way undermines Haack's point.

they are to be warranted, must ultimately rest on immediately warranted insights.[154] Thus, with coherence theories of warrant ruled out, the alternative to seeing intuition as the relevant warranting process seems to be mathematical skepticism.[155]

There are, then, several reasons for believing in intuition.[156] However, on their own, these won't be enough to convince intuition skeptics. This, at least in part, is because these reasons don't tell us what intuition is, or how it works.

## 5.2

## Intuition: Faculty or Capacity?

In saying what intuition is, a vital point on which to insist is that positing a process is not the same thing as positing a faculty. I do not hold that there is a faculty of intuition, if by 'faculty' is meant some dedicated mental module. The Platonist ought not, and need not, saddle himself with the project of describing some mental module, distinct from any other. He ought not, because faculty psychology is, to put it mildly, controversial, Fodor's attempt[157] to breathe life back into it notwithstanding. The Platonist's claims are controversial enough in their own right, he ought not link their fate to that of another controversial thesis. There is moreover, (so far as I am aware) no evidence from empirical psychology to support the claim that there is a dedicated

[154] I remind the reader that I am here considering only warrants which are internally accessible.

[155] See Katz (1998, 43-4) for a fuller discussion of this kind of argument "by elimination of alternatives" for the existence of intuition.

[156] These reasons for the belief that intuition is the process that produces our warranted (basic) mathematical beliefs also tell against Casullo's claim that one must warrantedly believe his (I) in order to have warranted intuitions. See section 4.1.

module for basic mathematical beliefs. Insisting on the existence of an independent faculty leaves the Platonist with the dilemma of either clinging to belief in such a faculty in the absence of any evidence, or moving to an indefensible mind/body dualism.

So, the Platonist ought not see intuition as a separate module. He need not, because there is another way to account for intuition. Katz (1998 and forthcoming) and Bonjour (1998) have recently offered accounts of intuition according to which intuition is simply an aspect of the ability to reason. Bonjour (1998, 109), holding that the claim that intuition involves a separate faculty is anything but obvious, writes that "the faculty involved is simply the ability to understand and think." Katz sees reason and intuition as part of the same faculty. His account, moreover, provides a way to do this while at the same time preserving their distinct roles:

> We can conceive of reason as our rational faculty in application to deductive structures and intuition as the same faculty in application to elements of such structures. Intuition is reason in the structurally degenerate case.[158]

One might argue that this move actually undermines its own analogy. One might, that is, argue that reason just is an ability to apprehend structural relations, and deductions. Take away these, and there's nothing "like reason" for intuition to be. But this would be a mistake. Take a basic mathematical statement like 'Four is composite'.

---

[157] Jerry Fodor, Modularity of Mind: An Essay on Faculty Psychology (Cambridge, Mass.: MIT Press, 1983).

[158] Jerrold Katz, "From the Philosophy of Mathematics to the Philosophy of Philosophy," Journal of Philosophy (forthcoming).

Although this would be an element in a deductive structure, this is not to say that there aren't necessary connections between the concepts used in the statement, which intuition can comprehend, as reason comprehends the necessary connections between premisses and conclusion in the deductive structure. When one comprehends the meaning of the word 'four', and that of 'composite', one sees that it could not be the case that four was not composite. Similarly with 'Three is an odd number'.

Intuition, thus far is the faculty of reason in the structurally degenerate case. We have yet to say anything about how it works. This is the task for the next section. In the course of completing this task, we will also gain a fuller picture of what intuition is.

## 5.3

## How Does it Work?

Supposing that we do in fact have intuition, a natural question to ask is how it works. This question can be interpreted in two different ways, however. On one reading, it asks how intuition provides warrant for basic mathematical beliefs; on the other, it asks a nuts-and-bolts question of the mechanics of the process of intuition. Taking the second question first, the situation is as follows. Given the way in which we have linked intuition and reason, it is safe to say that intuition works the way reason does. The precise way that mental or brain processes work when we form basic mathematical beliefs is a matter for investigation by empirical psychology and the neural sciences, just as is the precise way in which our mental or brain processes work when we use reason to problem-solve. In other words, this form of the question is not one for a priori philosophy.

A caveat. There is one issue here which might reasonably be pursued a priori, however, and that is the question of whether our ability to form beliefs about abstract objects must involve a dualism of the mind. Intuition, for Bonjour, involves universals literally being part of the content of one's thoughts. Bonjour recognizes that this is incompatible with a materialist psychology, but he regards this as reason to reject the materialist psychology. Having linked intuition with thought in general, Bonjour (1998, 181) writes:

> What is needed is an intelligible account of the connection between the intrinsic features of thoughts or thought elements and such things as metaphysically independent properties or universals, in virtue of which the former can be about the latter in an internally accessible way.

Bonjour goes on to offer an account which has its roots in Aquinas and Aristotle. For Bonjour, when I think of some X, the form of X literally occurs in my mind. To use his example, my mind is informed by the form of triangularity when I think of a triangular figure (Bonjour, 1998, 183). Since he is aware that triangularity itself does not appear in my mind when I think thoughts involving triangles, Bonjour holds that what does occur in my mind is a separate universal: triangularity qua having esse intentionale, as opposed to triangularity qua having esse naturale. Bonjour however, does not expand on what is meant by 'esse intentionale', nor does he have any account of how these different universals are related. Nor does he seem at all worried by the fact that his ontology now involves three modes of existence: physical states, universals (in the extra-mental sense), and universals in the (unexplained) "intentionale" sense.

What gets Bonjour into these troubles is his suspicion (1998, 181) that materialistic theories of mind cannot account for intuition (or, for that matter, for thinking), together with his conviction that materialism is far less obvious than the fact that we can think contentfully and be aware of those thoughts. While this latter point is true enough, it seems the more relevant comparison is between the obviousness of materialism and that of the account with which Bonjour seeks to replace it. And here materialism seems to fare better. For me, at any rate, if Platonist epistemology conflicts with materialist psychology, it is the epistemology which would have to go. I therefore reject Bonjour's account.

Fortunately, however, there is no reason to think that there is necessarily a clash between Platonist epistemology and materialist psychology. This will emerge from the following answer to the second form of the question under discussion. How does intuition provide warrant for our basic mathematical beliefs?

According to Katz (1998, 44)[159] intuition is a direct apprehension of the structure of an abstract object. These apprehensions of structure "reveal the limits of possibility with respect to the abstract objects having the structure" (Katz, 1998, 45). That is, intuition rules out the possibility of objects with a given structure having, or not having, a certain mathematical property. A mathematical intuition, then, is a non-inferential mental representation of abstract mathematical reality. The relationship between mathematical reality and minds is that of a representation which preserves both

[159] Katz's discussion is couched in terms of the justification of mathematical beliefs, rather than the warrant thereof. However, since he (1998, 36) cashes this out in terms of the grounds a believer has for holding his beliefs, and since I am here considering warrant which is available to the subject, it does not seem to me that there is any great gulf between us over the internally accessible cases. In any event, I am here transposing his account from the concept of justification to that of warrant.

structural relations between elements of that reality, and the mathematical properties of individual elements. Any system of representations which successfully preserved such properties and relations would be reliable, and thus warranting. Moreover, since this relationship cannot involve any causal component, it thereby cannot involve any essential use of experience. The process of intuition must therefore be an a priori process.[160]

For any given mathematical belief, the question of its truth is settled by whether it accurately describes mathematical reality. The question of its warrant depends, as reliabilism has it, on the reliability of the producing process. For basic mathematical beliefs, that process is intuition. As emphasized elsewhere in this dissertation, intuition is thought to be fallible, as any human capacity. This, however, does not prevent it from being reliable. That intuition is reliable is shown as follows. Firstly, the Platonist holds, what seems in any case obvious, that we have a large body of mathematical knowledge. Secondly, as was argued in section 5.1, there are good reasons for holding that we have intuition. Thirdly, as was also argued in that section, in order to account for our mathematical knowledge, we must have some warranted basic mathematical beliefs, and intuition is the only candidate for the process which produces these. Thus, we have mathematical knowledge, we have intuition, and intuition is responsible for that knowledge. Finally, given the volume of mathematical knowledge, and its certainty relative to other kinds of knowledge, mathematical intuition is highly reliable.

---

[160] As we shall see later in this chapter, there is some role for experience to play in intuition. But this is simply that of providing the concepts which will be the inputs for the process of intuition. The present account, then, is Kantian in this respect. It nevertheless preserves the apriority of intuition, since the process makes no use of experience once the concepts have been thus supplied.

The claims that we have intuition, and that it is responsible for our mathematical knowledge are, quite clearly, the controversial premisses in the above argument. However, let us attempt to remove some of that controversy by reminding the reader what I mean by these claims. Given that I am committed to a physical instantiation of intuition in the brain, and that this intuition is simply our ability to reason, to say that we have intuition is just to say that our minds are, once furnished with the relevant concepts, ordered in such a way as to allow us to understand basic mathematical truths. To say that intuition, in this sense, is responsible for our basic mathematical beliefs is just to say that our ability to reason enables us to establish the truth of basic mathematical beliefs, and produces in us belief in these truths—without any causal or experiential connection between minds and abstract objects.[161] The reliability of intuition is intimated, as before, by the volume and certainty of mathematical knowledge. Since, moreover, the process of intuition requires only the ability to form representations of mathematical reality, there is no obstacle to its repeatability. Intuition, then, is both reliable and repeatable. It is therefore, by reliabilist lights, a warranting process. Since it makes no use of any particular characteristic of experience,

[161] Mark Balaguer, in his Platonism and Anti-Platonism in Mathematics (New York: Oxford University Press, 1998) also attempts to reconcile Platonism, mathematical knowledge and standard semantics. His idea, which he calls Full-Blooded Platonism (FBP), has it that mathematical reality is so replete that any consistent theory describes some portion or other of it, without requiring any causal contact. While the position is a tempting one, it is also problematic. For example, since any consistent description describes some portion of the mathematical realm, there will be an infinite number of distinct mathematical universes. The axiom of choice (e.g.) will be true in some of these universes, but false in others. This, Balaguer (44) notes, has the consequence that mathematical truths are not necessary truths, and this is not something I can accept, given my use of necessity in previous chapters. Moreover, because there are infinitely many mathematical universes, there will be no such thing as (e.g.) the axiom of choice; there will only be the axiom of choice for a given universe (Balaguer, 59). Standard models are standard only because they are the "intended interpretations. This, coupled with Balaguer's (89) remark that standard semantics is "false but useful" seems to me to give up standard semantics, rather than capture it.

and in any case cannot make use of any causal connection between knower and known, it is not only a warranting process, it is an a priori warranting process.

Thus far, however, the account speaks only to the warrant attaching to mathematical beliefs, not to the accessibility of that warrant. How, on the present picture, does the subject know that he is warranted?[162] In answering this question, we must bear in mind the distinction between a belief being warranted, and a subject being warranted in believing. For the first, all that is required is that the belief be formed by a process which reliably and repeatably represents mathematical reality. For the second, in addition to this, the subject must have good reason to believe that his belief is of this kind. As we shall see in what follows, the methods by which one has access to one's warrant are themselves reliable belief-producing and sustaining processes, and thus are capable of adding to or detracting from the warrant deriving from the process described above. The point I would insist upon, however, is that the following processes are not necessary conditions on warrant. A Ramanujan-like figure can have warranted beliefs without having access to that warrant, though the warrant for his belief might increase if he were to go through these steps (his believer-justification would certainly increase).

The first way in which a subject can be warranted in believing some deliverance of intuition is the obviousness of basic mathematical truths. The proposition that 3 is an odd number, for example, is a truth so obvious that one would be at least prima facie warranted in believing it.[163] Understanding of the concepts involved is all that is needed

[162] Those who hold that, in order for someone to know something, he must know that he knows it, will, perforce, hold that some such story as the one about to be told, is in fact necessary for a subject to be warranted, not only for him to know that he is.

[163] Michael Devitt, commenting on an earlier draft of this dissertation, objected that it was hard to see how obviousness conferred warrant, since most of humanity thinks theism is obvious, and some large

to see the truth of the proposition. Since, ex hypothesi, this belief is produced by a process of intuition, one is thus at least prima facie warranted in believing the deliverances of intuition. This prima facie warrant can be increased or diminished by checking the relevant intuition by either of two methods. Firstly, the intuition may be checked against other intuitions (Katz, 1998, 46). If it is consistent with these, the subject's initial warrant increases. If it is not consistent with them, then his warrant for at least one of the intuitions is destroyed. It may be obvious which of the intuitions is causing the problem. If so, then that intuition is the one to reject. In the event that it is not so obvious which intuition is to be rejected, appeal can be made to the second method.

The second method by which one can check on whether one is warranted in believing a given intuition is its fit with relevant sections of the corpus of accepted mathematics (Katz, 1998, 46). If it is inconsistent with accepted mathematics, the subject's intuition loses its prima facie warrant. If it is consistent with accepted mathematics, his warrant increases. The warranting of the deliverances of intuition thus takes the form of a reflective equilibrium model.

In the event, discussed two paragraphs back, that we are attempting to decide which of two conflicting intuitions is warranted, we would check both against relevant

---

proportion of them may think that the doctrines of transubstantiation and the Trinity are obvious. The objection fails, however, because we are concerned here with beliefs which can be seen immediately to be obviously true, and whose warrant does not require support from other beliefs. None of the examples Devitt notes fit this condition; even the most devout have their moments of doubt, and, more importantly, have reasons for their belief in god (most often, some rudimentary version of the argument from design). Their beliefs then, are neither obviously true to them, nor immediately warranted. Furthermore, obviousness only supplies an initial degree of warrant. This initial warrant must then be tested for consistency against other intuitions (see below). Even if the beliefs listed by Devitt were warranted to some degree because of their obviousness, this warrant would not survive the kinds of tests outlined below.

portions of accepted mathematics. At least one of the conflicting intuitions will be inconsistent with portions of accepted mathematics, and thus would be rendered unwarranted. Upon seeing this, the subject would no longer be warranted in believing the conflicting belief.[164]

There is, of course, the possibility that currently accepted mathematics is inconsistent, and thus that each of two inconsistent intuitions could appear to be consistent with it.[165] There are two reasons for thinking that this possibility does not represent a serious worry for the epistemology outlined here. Firstly, it seems likely that if accepted mathematics were inconsistent, somebody would have noticed this by now. The fact that no such discovery has been made, in a discipline where every piece of reasoning is checked and rechecked by many mathematicians, and where such a discovery would make someone's career, surely gives us reason to think that accepted mathematics is in fact consistent. Secondly, if accepted mathematics were inconsistent, then the reflective equilibrium method offered here would eventually bring this to light. If accepted mathematics were inconsistent, then, at some point, when intuitions

---

[164] Once we admit the distinction between a belief being warranted, and a subject being warranted in believing it, some interesting cases open up. For example, it might be the case that both intuitions were inconsistent with accepted mathematics, and that, after investigation, the subject were to discover this. In such a case, of course, the subject, in believing either of these, would not be accurately representing mathematical reality, and neither belief would be warranted. In discovering this, the subject would become unwarranted in believing either of these. Or, it might be the case that neither of the inconsistent beliefs produced by intuition were known to be inconsistent with the rest of accepted mathematical knowledge. Since at least one of them must be, at least one of these beliefs would not be warranted. However, the subject, ex hypothesi, would not know which one was the unwarranted one. Since he would know, however, that at least one of the conflicting beliefs must be unwarranted, he would not be warranted in fully believing either. The only thing the subject would be warranted in doing in such a case would be withholding belief in both, at least for the time being. The possibility of such cases, however, presents no objection to seeing intuition as an a priori warrant-conferring process. At worst, it shows that intuition is fallible.

[165] Thanks to Michael Levin for raising this point.

concerning accepted mathematics were being checked against each other, the inconsistency between two propositions of accepted mathematics would show up. At this point, one of the propositions causing the inconsistency would be rejected. If the correct choice of which proposition to reject were made, mathematics would then become consistent. If the incorrect choice were made, other inconsistencies would show up later, until the problematic proposition was rooted out.

There is, then, no obstacle to a given subject having access to the a priori warrant attaching to beliefs produced by intuition. Of course, it may not always be easy to access in this way the warrant attaching to a belief produced by intuition, but this point has no epistemic significance.

The above account shows how warrant might accrue to basic mathematical beliefs. However, one might object that it requires that subjects have a type of Cartesian access to their (basic mathematical) concepts, and that there is no good reason to believe that subjects have any such access.[166] For us to have Cartesian access to our concepts, the mind would have to be transparent to itself, and allow us, infallibly, to know our concepts. If Katz's account required this kind of access to concepts, then, I agree, it would be dead in the water. But there is no reason to think that it requires this kind of access to concepts. All it requires is that we have reliable access to our concepts in enough detail that so that we can have a good—but fallible—idea of the content of our thoughts. This kind of access we surely have. And, once we do have such access, and can form some beliefs about our mathematical concepts, then the reflective

[166] Michael Devitt raised this objection in connection with an earlier draft of this dissertation.

equilibrium model outlined above will correct for any initial errors that may arise due to our having imperfect access to our concepts.

As the processes by which one can access one's a priori warrant are themselves a priori mathematical processes, they would themselves be sufficient to provide a priori warrant. As the process of intuition, unaided by such processes, is also an a priori process, these other processes are not necessary conditions on warranted mathematical belief. Bare intuition, unchecked, is a warranting process. Checking one's intuition via these other processes is not only warranting, but epistemically responsible. Checking the warrant provided by intuition in this way helps one to see whether the defeasible warrant provided by intuition is, or is not, defeated. It also provides a check on whether the warranted belief provided by intuition is both warranted and true, or merely warranted-but-false.

To sum up, mathematical intuition is the fallible faculty of a priori reason, directly representing the elements of the deductive structure of mathematical reality mathematical intuitions are warranted insofar as they reliably correspond to mathematical reality, and we can have confidence in the truth of any given intuition, p, on the basis of the reliability of intuition, the systematic fit between p and other intuitions, and the systematic fit between p and accepted mathematics. Since going through these checks and balances is itself a reliable process, this process is capable of increasing or decreasing the amount of warrant attaching to a belief due to its correspondence with mathematical reality.

The foregoing raises several interesting questions. Even on the assumption that we in fact have intuition, the following require answers.

Q.1. How does the representation of mathematical reality required by the above account come about? I.e., even if we do have intuition, and there is good evidence for that, how does it come to pass that physical creatures can have representations of abstract objects?

Q.2. Could the process, as described, yield a priori warrant?[167]

Q.3. How is the Field-Maddy requirement that we explain the reliability of our mathematical beliefs satisfied?

Q.4. How, given that we have access only to the mental side of the relationship between our representation of mathematical reality and that reality itself, can we know that a correspondence holds between representation and reality?

I shall examine each of these questions in turn.

## 5.4

## How Might the Required System of Representations Originate?

We have seen that, in the absence of any possibility of causal contact between ourselves and a realm of abstracta, the warrant possessed by our mathematical beliefs arises due to the reliability of the process of forming representations which correspond to

[167] By the criteria outlined earlier in this chapter (repeatability and reliability) the answer to this question would appear obviously to be positive. Kitcher, however, has an important objection to the possibility of any form of intuition yielding a priori warrant, which must be dealt with before we can claim to have fully answered this question.

mathematical reality itself. This raises the question of how it is that we might form such representations. To answer this question, we must bear in mind that Platonists typically draw a distinction between the genesis of our mathematical beliefs, and their warrant. The Platonist story told here about the a priori warrant of mathematical beliefs has to do with a reliable correspondence holding between mathematical representation and reality. When it comes to explaining the genesis of our mathematical beliefs, the Platonist, however, can help himself to the same story typically told by nominalists and other anti-Platonists. The story is a familiar one. In order to keep track of who owned which livestock, our forefathers developed rudimentary systems of counting. In order to keep track of whose tract of land ended where, and in order to facilitate the construction of buildings, they developed similarly rudimentary systems of geometry and trigonometry.

Since, on the Platonist conception of mathematics, mathematical truths are necessary, and mathematical objects have their intrinsic and structural properties necessarily, it follows that instantiations of mathematical objects or properties in the physical world will preserve the intrinsic and structural properties and relations pertaining to the abstract originals. Thus, triangular shaped objects will necessarily have three sides, and the length of the circumference of a circular shaped object will necessarily be double the product of pi times the radius. The general point here is similar in one respect to one of the main claims made by mathematical structuralists. The physical world contains instantiations of mathematical patterns, and our beliefs about mathematics arise initially through causal interaction between us and those concrete instantiations.[168] The difference between my position and that of the structuralist is that I do not go the further step and identify the numbers with positions in patterns. The causal interaction between us and the concrete instantiations here serves only to furnish the mind with mathematical concepts. Our mathematical beliefs become warranted when intuition enables us to understand the truth or falsity of propositions containing those concepts:

On this picture, then, the mathematical beliefs of our ancestors reliably corresponded to (rudimentary aspects of) mathematical reality, because the concepts constituting the content of these beliefs initially arose out of physical interactions (e.g., perception) with physical instantiations of that mathematical reality. Since the objects thus interacted with preserved the mathematical properties of the abstract originals, interacting with such objects allowed our mathematical ancestors reliably to form representations of mathematical reality—and thus to develop rudimentary systems of mathematics. These rudimentary systems have been refined and extended by each succeeding generation, allowing us to arrive at the rich, sophisticated mathematics we have today. These more sophisticated developments result in more detailed representations, corresponding to more complex aspects of mathematical reality. The reliability of these more complex representations is not explained solely in terms of

[168] To avoid any potential confusion, I should also distinguish my position from that of Maddy (1990). Maddy held that we came to have knowledge of mathematical objects by actually physically interacting with the objects themselves: thus, I come to have knowledge of the number 6 by interacting with the half-dozen eggs in my fridge (and other physical groups having six elements). On this position, the number 6 is itself instantiated wherever there is a physical entity having six parts. My position has it instead that we do not interact with mathematical entities, since these are abstract. Rather, we interact only with physical objects which instantiate certain mathematical properties. We thus get our initial knowledge of mathematical objects through interacting with physical objects that are, as it were, at one remove from the mathematical objects themselves, but which preserve enough of the mathematical properties of the originals so as to enable us to form accurate representations of the originals.

interactions with physical instantiations of abstract mathematical objects. Rather it is explained by the checks and balances imposed on any putative addition to our body of mathematical beliefs. Any such belief must at least not be in contradiction with accepted mathematics, must show itself to be fruitful, and must be rigorously examined by many trained mathematicians. We can have confidence in the claim that our representations of mathematical reality accurately correspond to that reality to the extent that we can have confidence in the reliability of such checks and balances.

The process of building more complex representations, as shall be emphasized in the following section, is, like all human enterprises, thoroughly fallible. For now, it is enough to note that the apriorist Platonist can explain both the genesis and warranting of our mathematical beliefs in unobjectionable ways. The warranting is occurs via a reliable, though non-causal-contact, process of forming representations of abstract reality; the representational side of that correspondence arises through causal contact between human minds and physical structures which instantiate features of mathematical reality.

A final point: my use of the distinction between genesis and warrant, taken together with my overall account of a priori mathematical warrant, should not be taken as a denial of the possibility of empirically warranted mathematical beliefs. Mathematical beliefs can be warranted empirically. A child's belief that 2+2=4, based on an induction from observed instances would be so warranted. Empirical warrant, however, though it might be enough for everyday purposes, does not yield certainty. Because of this, it would not suffice for the kind of knowledge sought by mathematicians; it would not suffice to establish mathematical truth. Nor, for the same reason, could empirical warrant outweigh a priori warrant.[169]

## 5.5

## Can Intuition Yield A Priori Warrant?

According to Katz's account, the warranting of mathematical beliefs is a matter of those beliefs reliably corresponding to the structure of abstract mathematical reality. According to the line I have sketched in the previous sections of this chapter, this correspondence arises out of two factors: our contact with physical objects which instantiate mathematical structures, and our reflection on the concepts so gained. Since intuition is both reliable and repeatable, it is a warranting process; since it involves no causal contact between us and mathematical reality, and thus allows no room for warrant which varies with experience, it is an a priori warranting process. Our access to the warrant provided by intuition is due to the obviousness of basic mathematical truths, and to our ability to check the fit between, on the one hand, a given intuition, and on the other hand, other intuitions and accepted mathematics.

Kitcher, though not directly addressing Katz's account, argues that no process of intuition could yield a priori warrant. In order to make his case, Kitcher introduces the notion of a suspect process. A process is suspect if a person could carry out this process, but would be aware of his inability to discriminate the type of process performed from other processes known to generate false beliefs (Kitcher, 1983, 62). Kitcher (1983, 63)

[169] To put the point in the language of Chapter Three, empirically warranting processes do not reliably track the truth of necessary truths, whereas a priori warranting processes do.

then claims that in order to show that intuition cannot yield a priori warrant, it would suffice to show both that Platonist intuition is a suspect process, and that there are experiences which can suggest the falsity of mathematical statements.

And, according to him, the first of these points can be easily shown by examples of our mathematical ancestors having hailed some principle as intuitively evident, only for that principle to turn out to be false. For example, Kitcher (1983, 63) describes both Frege and Cantor as having thought that it was "intuitively evident" that any property determines a set. The idea then is that, since we cannot discriminate intuitions from processes known after the fact to produce false beliefs, intuition is suspect. Kitcher (2000, 75) makes the same kind of point in a more recent article, where he writes:

> Appeals to elusive processes of a priori reason ought always to be accompanied by doubts about whether one has carried out the process correctly, and whether, in this instance, the deliverances are true.

But neither Kitcher's examples, nor the possibility of error inherent in any given deliverance of intuition suffice to show intuition to be suspect, and this for several reasons.

Firstly, since we have jettisoned Kitcher's condition (c),[170] we are under no obligation to hold that intuition is infallible. We can consistently hold that Frege's

---

[170] For the reader's convenience, I here repeat Kitcher's (1983, 24) conditions on a priori warrant:
W is an a priori warrant for X's belief that p iff W is a process such that, given any life e, sufficient for X for p,
(a) some process of the same type could produce in X a belief that p,
(b)if a process of the same type were to produce in X a belief that p, then it would warrant X in believing that p,
(c) if a process of the same type were to produce in X a belief that p, then p.

intuition was responsible for the relevant false belief,[171] and that intuition is a generally reliable, a priori process. A fallibilistic process of intuition is not shown to be suspect by pointing out one or two high-profile errors for which it may have been responsible. It would have to be shown that intuition was unreliable over a significant percentage of cases. And, it is very difficult to imagine how this could be shown for an understanding of intuition according to which it is responsible for our most basic, non-controversial mathematical beliefs.

Secondly, to depart for a moment from consciously accessible warrant, a process is not rendered suspect by the subject's being aware of his inability to tell that type of process from one known to generate false beliefs. It is in any case rather odd that Kitcher, who explicitly rejects the requirement that warrant be accessible to introspection, should think that it would. For a reliabilist, and Kitcher counts himself as one, once the belief is in fact produced by a reliable process it is warranted. This point should hold regardless of whether the subject has doubts about which process was responsible for that belief. That is, if intuition is genuinely reliable, and is responsible for our mathematical beliefs, then, beliefs produced by intuition are in fact warranted, even if the subject has erroneous second-order doubts about the reliability of the producing process.[172]

What then about the subject's access to his warrant? Surely one who could not tell whether his belief was produced by intuition, or by some unreliable process, would

[171] Actually, given our conception of intuition as a process which produces basic, immediately warranted beliefs, it is not likely that intuition was responsible for the kind of belief Kitcher is discussing. I relegate this point to a footnote, however, since Kitcher's point does not depend upon the particular examples.

[172] The case of Blinky, described in section 1.3.1, shows this, or so I claim.

not know that he was warranted. This is true, but harmless to the epistemology outlined above. If the subject has his doubts about the veracity of a particular mathematical belief, he can check that belief against relevant portions of accepted mathematics, and against other intuitions. There are, in other words, checks and balances in the reflective-equilibrium epistemology outlined above, which allow a subject to determine whether, and to what extent, a given belief is warranted. He is not, as Kitcher's picture might suggest, stuck in the hapless position of not knowing how his belief is produced, and not having any other way to tell whether his belief is produced by a suspect process, or is in fact warranted.

Thirdly, we can ask whether warrant is supposed to be an all or nothing thing. Jettisoning (c) leaves us with no reason to suppose that it must be, and so the following response is open to us. Intuition, unaided, yields a certain degree of warrant, n. But any deliverance of intuition must be integrated with (at least most of) accepted mathematics. Such theoretical desirables as fruitfulness, freedom from contradiction, systematic elegance, and so on, then have to be taken into account. The process(es) of checking the deliverance of intuition against these theoretical desiderata can add to, or subtract from, the degree of warrant conferred by intuition on the belief. This, assuming intuition was responsible for the relevant belief, would appear to be exactly what happened in Frege's case. The initial axiom was intuited; thus it was produced by a reliable, a priori process, and thus it had some degree of a priori warrant. Further investigation, however, led to the paradoxes, and so the initial warrant was undermined.

Hence, showing that intuition can lead to false beliefs would not be enough to show that intuition is a suspect process. Intuition can yield false beliefs, and those beliefs can be warranted, initially, to some degree. What this shows is that intuition is fallible, and so its pronouncements need always to be checked in other ways. It does not show that it is "suspect" in any damaging sense.

Fourthly, on the assumption that we do have a substantial body of mathematical knowledge, intuition clearly yields true mathematical beliefs in the overwhelming majority of cases. Although there are many examples of controversy in the history of mathematics (as Kitcher documents), the vast majority of mathematics is uncontroversial. Thus, most mathematics is produced by processes that are clearly reliable. Where the mathematical beliefs in question are suitably basic, the producing process is intuition. It is enough, in order for a process not to be suspect, that it be broadly reliable. Infallibility is not required. Given that many of our mathematical beliefs are true, intuition satisfies this requirement.

Fifthly, the errors Kitcher notes were discovered by using a priori reasoning. Russell's paradox is, if anything is, an a priori argument. Thus, once the infallibility requirement on intuition is rejected, Frege's and Cantor's misleading intuitions have no tendency to show that intuition is not a priori. The whole process can be seen as fallible, a priori intuition correcting itself (see Chapter Three above for more on this). It is part of the fallibilist position that an a priori warranted belief is open to having its warrant undermined by other a priori processes. The case does not show intuition to be unreliable, it simply shows that we must take care in checking the beliefs produced by this process. Nor is this response dependent on Kitcher's particular examples. Any case where intuition at first provides a subject with a false belief, but where that subject, on the basis of further examination of the consequences of accepting that belief, later rejects it, would appear to be a case of fallible, a priori intuition correcting itself.

Recall that Kitcher needed to show both that intuition was suspect, and that there could be experiences which suggested the falsity of mathematical statements. I have just argued that he has not shown the first, but nor has he shown the second. Kitcher (1983, 64) claims that mathematical beliefs can be undermined by one type of experience:

> [T]here is at least one type of experience which can suggest to us that the beliefs we have formed on the basis of intuition are incorrect. Mathematical beliefs are vulnerable to social challenges. Such challenges pose a sufficient threat to make it unreasonable for us to form beliefs on the basis of intuition when we recognize that we cannot discriminate our intuitions from processes which misled our predecessors.

A social challenge to a mathematical belief, as distinct from a direct challenge, takes the following form. Smith's mathematical belief in the truth of p, a belief warranted by the fact that it was produced by his mathematical intuition, is challenged by Jones. The challenge, however, is not along the lines of presenting an alleged refutation of p (this would count as a direct challenge). Instead, Jones is an expert, and on the grounds of his expertise, (and perhaps that of other experts who he claims are in agreement with him), challenges Smith's belief in p. Kitcher's claim, then, is that in cases such as these, we have experience suggesting the falsity of mathematical beliefs, and thus undermining their warrant. Granted this description of such cases, mathematical beliefs would not count as a priori.

Kitcher is unquestionably correct that there can be kinds of "experience" which, in some way, "suggest" the falsity of some mathematical belief which has been produced by intuition. To the kind of example offered here, we might add the following experiences: dreaming that some basic mathematical belief was actually false, ingesting LSD and having it appear to one that some basic mathematical belief was actually false, and being hypnotized to believe that some such belief was actually false. That this list might be continued indefinitely serves to show that while suggestions are cheap, not all need to be seriously considered. Thus, if suggestions are to have the wideswingeing epistemological consequences that Kitcher is after, we will have to make a distinction between experiences which seriously suggest the falsity of certain basic mathematical beliefs, and cheap suggestions such as those illustrated above.

Now, suppose we have narrowed down the class of "experiences" which suggest falsity so that it includes only serious suggestions. There is a further question to be asked at this point. Is the "experience" in question accurately described as an experience, or is it an experience in the broader sense in which any mental event in a subject's history, including that of understanding a proof, counts as experience?[173] Take the kind of case Kitcher has in mind, that of the devious mathematical expert, out to (falsely) convince you that one of your basic mathematical beliefs is false. If this social challenge is genuinely to count as an experience, it cannot be accompanied by a (misleading) refutation of your belief. If it were so accompanied, the case would be that of one set of a priori considerations being taken (mistakenly as it happens) to undermine another, and would present no problem for the apriorist. I will refer to experiences which are unaccompanied by any a priori reasons as "bare experiences." Thus the case

[173] See Chapter Three for a discussion of these two notions of experience. For now, however, it is sufficient to note that only the undermining of warrant by experience in the narrower (a posteriori) sense would be damaging to the apriorist's claims.

must be one where (i) one is presented with a bare experience of a mathematician telling one that one's belief is false, and (ii) this experience constitutes a suggestion serious enough to undermine the a priori warrant possessed by the subject for his belief (if the subject has no warrant, then of course there is no counterexample to apriorism). But (ii) is precisely what has not been shown by Kitcher. Kitcher assumes that the bare experience of having a mathematician tell you that your belief is false is enough to undermine the a priori warrant possessed by the belief, but he does not argue for this claim. He thus begs the question against apriorism.

One might attempt, along the following lines, to bolster Kitcher's example in order to show that such an experience would undermine the subject's a priori warrant for his belief. Suppose that the subject in question told the mathematician that he had a priori warrant for his belief, and that the mere suggestion that it was false was not sufficiently serious enough to render his belief unwarranted. If the mathematician wanted to pursue his evil ways, he would have to offer some reason for thinking that the belief in question were false. That is, he would have to offer intellectual considerations which appealed to the understanding of the subject. But if he were to do this, the case would thereby move beyond that of a bare experience, and (i) is no longer satisfied by the case. The case would thus not be one where one had an experience which suggested the falsity of a mathematical belief.[174]

The key point here is that a social challenge will not be strong enough to undermine the a priori warrant conferred on some p by intuition. For all the social

[174] In Kitcher's terminology, we no longer have a social challenge, but have instead a direct challenge. The threat posed to apriorism by such direct challenges was met in Chapter Three.

challenge can yield is the promise of a reason to reject p. If this is all that is forthcoming, then one would have only the mere possibility that intuition is misleading you. Such a mere possibility does not justify the epistemological radicalism advocated by Kitcher. If, on the other hand, more is offered, then the case is no longer a social challenge, but a direct one, and Kitcher has not shown that there are experiences which seriously suggest the falsity of mathematical beliefs.

Of course, if intuition were genuinely suspect, in any damaging way (i.e., if 'suspect' amounts to more than simply 'fallible'), then it might well seem that a direct challenge of the sort outlined might serve to undermine a subject's warrant for his basic mathematical belief. Kitcher, taking himself to have shown that intuition is suspect, is of the opinion that he has done enough to show that the possibility of social challenges defeats apriorism. However, since I have offered reasons which show that intuition has not been shown to be suspect, but has rather been shown merely to be fallible, I see no reason to see social challenges as a serious threat to apriorism.

However, even granted that intuition might be suspect, the existence of social challenges does not refute apriorism. Such challenges might suffice to deter mathematical unsophisticates from their mathematical beliefs, but so long as there are mathematicians whose a priori warrant for their beliefs outweighs the epistemological threat posed by social challenges, the apriorist need lose little sleep over such challenges. And surely no mathematician would just take the word of another that some basic mathematical truth was in fact false. They would, quite warrantedly, continue on in their beliefs unless given an actual reason for rejecting the belief in question. And, so long as there are some subjects who have a priori warrant for their beliefs, the apriorist wins the epistemological debate.[175]

Kitcher, then, had attempted to show that intuition, however conceived, could not confer a priori warrant on mathematical beliefs. This, he held, was because the combination of (a) the suspectness of intuition, and (b) the existence of social challenges, was sufficient to refute the apriorist's claims about the warranting power of intuition. If the case had been made successfully, then of course the Katzian account of intuition I am defending would have been refuted. The case has not been successfully made, however, and I conclude that Kitcher has offered no compelling reasons to think that intuition cannot confer warrant on basic mathematical beliefs.

## 5.6

## Satisfying the Field-Maddy Requirement

Recall the Platonist assumption that mathematical objects have all of their intrinsic and structural properties necessarily. As we have seen, this means that all physical instantiations of mathematical objects will preserve these properties. Although we cannot interact with the abstract originals, we can interact with the physical instantiations of these, and in so doing, we obtain our basic mathematical concepts; these concepts are then available to reason. We can see, for example, by examining our concept of a triangle, that it is impossible for an object to be triangular if it did not have

[175] In fact, even a successful wholesale deception of the entire mathematical community would not suffice to refute apriorism. One might define an a priori truth as one which could be known through purely a priori methods. If there are such truths, they would persist even though no-one actually knew them a priori. Mathematical truths would thus be a priori truths, knowable, even though not actually known, a priori.

three sides. Note at this point that the project of explaining the reliability of our beliefs about abstract objects has been reduced to the problem of explaining the reliability of our reasoning about empirically gained concepts. This problem is much less imposing than the Benacerraf problem. Moreover, it is not a problem that the anti-Platonist can use to cause trouble for the Platonist, since everybody agrees both that reason is reliable, and that it takes empirically obtained concepts as inputs. There is thus no special problem for the Platonist here: all sides to the dispute over abstract objects are ultimately obligated to provide some account of how it is that reasoning about empirically gained concepts is reliable. We have thus reduced an apparently intractable problem which arises directly out of the Platonist's metaphysics to a metaphysically neutral problem about how reason works.

One might object that the empiricist has an advantage here over the Platonist apriorist. That is, the empiricist can agree that reason is reliable, and that all of our concepts are obtained via experience, but go on to give an account of the reliability of reason in terms of the empirical testability of our beliefs. In other words, the empiricist holds that for any belief that has been produced by our reasoning abilities, we can design an empirical test for that belief. Once the belief in question has passed sufficiently rigorous testing, we can have confidence it its truth; once enough beliefs pass such testing, we can warrantedly believe that reasoning about empirically gained concepts is a reliable process. The Platonist, so the objection concludes, has no such test for his claim that reasoning about mathematical concepts is reliable, for there are no empirical tests available to test our mathematical beliefs.

The objection, however, fails. It is true enough that there are no empirical tests with which to verify our mathematical beliefs, but this is not to say that there are no tests. The process of reflective equilibrium, outlined above, is a method of testing our mathematical beliefs. Basic mathematical beliefs are checked for consistency with each other, and with the rest of accepted mathematics. Those beliefs which pass this testing process are accepted as true, and once enough beliefs pass such testing, we can warrantedly believe that reasoning about empirically gained mathematical concepts is a reliable process.

There is another reason for doubting that the Field-Maddy requirement has genuinely been solved. Consider reliable empirical processes; these can yield knowledge, even though they sometimes produce beliefs which are mistaken. Part of explaining the reliability of such processes is explaining what goes wrong in the instances where the reliable process produces belief in a falsehood; thus, part of explaining the reliability of vision is explaining what goes wrong in cases of (e.g.) hallucination. But there doesn't seem to be any way of explaining what goes wrong when reliable, but fallible, a priori processes produce beliefs in falsehoods. Thus, we have not genuinely explained the reliability of a priori mathematical processes.[176]

In fact, there are explanations available for what goes wrong when reliable a priori processes produce mistaken beliefs. Consider, once more, Frege's erroneous belief that every property determines a set. The principle that properties determine sets holds of many properties, but what applies to many properties does not necessarily

[176] My thanks to Michael Devitt, who raised this objection in connection with an earlier draft of this dissertation.

apply to all properties. And, in fact, the very nature of the paradox which showed that something had gone wrong showed what it was that had gone wrong: a principle had been applied beyond the proper boundaries of its application. Consider also another category of mistaken mathematical beliefs: those about what holds at limit cases. Such mistakes arose due to belief in the principle that whatever holds close to a limit must also hold at that limit. Reasoning processes that are reliable when applied to any point approaching a limit turn out not to be reliable at that limit. Mathematicians unaware of this used these processes for limit cases, and were thus led to believe falsehoods. Again, the very nature of the problem shows where our reasoning has gone wrong: reliable processes have been applied beyond their proper boundaries of application, and have thus had their reliability undermined. These cases seem entirely parallel to cases such as hallucination, where a reliable empirical process is applied under conditions which undermine its reliability. The explanation in both kinds of cases is the same: we explain what it is about the conditions in which the process led to belief in falsehood that render the process unreliable in such circumstances. There is, in general, then, no reason to believe that we cannot explain what goes wrong when reliable a priori processes lead to mistaken beliefs, and the objection fails.

## 5.7

## How Do We Know That Our Beliefs Correspond to Reality?

Suppose it is agreed that a correspondence between our representation of mathematical reality, and that reality itself, might be possible in the absence of causal contact between ourselves and mathematical reality. Suppose further that it is agreed that intuition could warrant our mathematical beliefs in something like the way here sketched. There remains a further question. How do we know that the required correspondence actually exists, as opposed to being merely possible?

Though the question is a genuine one, it is worth bearing in mind that it takes us beyond the Benacerraf problem, as applied to reliabilism. That problem was to show how reliabilism could account for our mathematical knowledge, on the assumption of a Platonist ontology. This question has been answered in the course of Chapters Three, Four, and in the preceding sections of the current chapter. Thus, the main work of this dissertation is complete. I shall, however, all too briefly, outline an answer to this question below.

There are two ways in which the question might be taken. The first takes the question to ask how we know of a particular belief that it corresponds to the relevant aspects of mathematical reality. The answer here is decided in accordance with the reflective equilibrium epistemology provided above. Of paramount importance in answering this question is how the belief fits with known results, the fruitfulness of accepting the belief, and the initial warrant accruing from the obviousness of the intuition.

The second way the question might be taken is more general. This form of the question asks how we know that a correspondence holds between our mathematical beliefs, taken as a whole, and mathematical reality. The Platonist's answer here is that, in general, we know that the correspondence holds, because we know that we have mathematical knowledge, we know that this is of a platonic realm, and we know that such knowledge could not be causal. Our beliefs therefore must correspond to mathematical reality. More formally:

I.1. We have mathematical knowledge (ass.).

I.2. This knowledge has a realm of abstract objects for its object (ass.).

I.3. Knowledge of abstract mathematical objects must be either causal or take the form of a non-causal relation of correspondence between mathematical reality and our representation of it (ass.).

I.4. Knowledge of abstract mathematical objects cannot be due to a causal process that links us with such objects. (Benacerraf problem).

I.5. Knowledge of abstract mathematical objects must arise from a non-causal correspondence between mathematical reality and our representation of it (D.S., I.3, I.4).

I.6. There must be a non-causal correspondence between mathematical reality and our representation of it. (M.P, I.1, I.5).

(I.1) and (I.4) are generally agreed. (I.3) seems unobjectionable. The only controversial element of the Platonist's answer to the general form of the question, then, is his (I.2)—the claim that our mathematical knowledge is knowledge of a platonic realm. The defense for this claim derives from the theoretical virtues of Platonism (outlined in the introduction to this chapter). The theoretical strengths of Platonism warrant belief in intuition, and Benacerrafian considerations show that such intuition can only work through a non-causally based correspondence between mathematical representation and reality. The short answer, then, to the general question, is that we know that there is a correspondence because our best theory of mathematics tells us there must be one. However, not only does the theoretical strength of Platonism warrant belief in the required correspondence, but also, in showing how an non-causal-contact intuition is possible, we refute the objection that Platonism is epistemologically bankrupt. That is, in showing the non-causal correspondence to be possible, we further strengthen the warrant possessed by Platonism. The relationship between Platonism and the non-causal correspondence it requires is thus mutually warranting.

If, in our answer to the general form of the question, we end up holding that we know that the right kind of correspondence exists because Platonism requires it, is there any improvement on the traditional argument for intuition—that there must be such a capacity, because Platonism requires it? Well, yes. Firstly, intuition has not been posited merely because Platonism requires it. Several reasons for believing in intuition were offered, and a description of the process was presented. The account given neither requires any "mysterious aphysical grasping" (or mysticism, for that matter), nor does it pin its hopes on the existence of a separate mental faculty. Thanks to the work of Bonjour and Katz, we can conceive of intuition as being simply one aspect of the ability to reason. Benacerraf has been answered on how mathematical knowledge is possible on the assumption of Platonism. We have, in the course of providing that answer, provided a reliabilist account of mathematical knowledge, and a reliabilist account of intuition. Further, we can explain how the required correspondence arises, in a manner that would seem to be perfectly acceptable to the empiricist anti-Platonist. It is a long story, but is ultimately to be explained in terms of the instantiation of mathematical properties in the physical world. We have also satisfied the Field-Maddy requirement that the reliability of intuition be explained. There is also one final consideration which supports the claim that we are not merely urging that there must be intuition because Platonism requires it. Since we have outlined and defended a Platonist epistemology, it is safe to say that the objections to the very possibility of there being mathematical intuition have been found lacking. Given all of this, what more could fairly be required of an answer to the Platonist's epistemological problem?

Epistemology was supposed to present an insurmountable challenge to the Platonist. We have seen that a very plausible account of the origin of mathematical beliefs gives us good reason for believing possible the required correspondence between mathematical beliefs and mathematical reality. And the reliability of that correspondence is what warrants those beliefs. The Platonist's insurmountable problem turns out to be fairly surmountable after all. Since, epistemology to one side, the Platonist has the best overall account of mathematics, and since his epistemology is not so hopeless after all, we should accept the Platonist's account. Along with the intuition-driven epistemology which accompanies it.

## 5.7

## Concluding Remarks

This dissertation has offered an extended answer to the Benacerraf problem, by attempting a reconciliation between a priori Platonism, and reliabilism, an epistemological theory usually thought to be both a posteriori and anti-Platonist.

The work of the first two chapters was preparatory. In Chapter One, I defended reliabilism against various objections, and found these objections unable to withstand scrutiny. In Chapter Two, the concept of a priori knowledge was defended against the Quinean attacks, which were likewise found wanting.

Finding no insurmountable objection to either reliabilism or the a priori, the next step was to put the two together, so as to yield an acceptable epistemology of mathematics. In order to be successful, such an epistemology was required to show the processes responsible for our mathematical knowledge to be a priori, reliable, and repeatable. Chapter Three argued that there were no sound objections to the traditional notion of mathematical proof as an a priori method. It furthermore found proof to be a reliable method, since the essence of an abstract mathematical proof is that it necessarily confers warrant on its conclusion. Since, moreover, formal proofs consist entirely of premisses either known to be correct, or assumed for the sake of argument, and of sanctioned transition rules, proofs are also repeatable. Thus, proofs were found to be reliable and a priori warranting processes.

Chapters Four and Five articulated and defended a concept of a priori mathematical intuition. Chapter Four argued that, contra Casullo, there were no serious obstacles to a reliabilist solution to the Benacerraf problem which made use of the notion of a priori mathematical intuition. The present chapter has offered a non-causal account of mathematical intuition according to which, intuition makes no essential use of any particular characteristic of experience. It has thereby shown intuition, thus conceived, to be an a priori process. Further, this process has been shown to be both repeatable and reliable. Moreover, the Field-Maddy requirement of explaining that reliability has been satisfied. Thus, it has been shown that the process of a priori intuition is, by reliabilist lights, a warranting process.

If what I have argued is, by and large, correct, then Benacerraf is answered by providing a reliabilist account of a priori mathematical knowledge which neither requires nor allows causal contact between mathematical reality and the knowing subject. If I am right, not only has the main problem for the apriorist Platonist been solved, but also one of the main objections to reliabilism—that it cannot account for our mathematical knowledge—has been refuted. That is to say, if what I have argued is correct, then not only has the apriorist Platonist an answer to Benacerraf, but the reliabilist is now able to account for mathematical knowledge, whatever the correct ontology of mathematics might be.

Bibliography

Alston, William. "How to Think About Reliability," Philosophical Topics (spring 1995): 1-29.

_____. "Level Confusion in Epistemology," Midwest Studies in Philosophy V (1980): 135-150.

Azzouni, Jody. Metaphysical Myths, Mathematical Practice: The Ontology and Epistemology of the Exact Sciences. Cambridge: Cambridge University Press, 1994.

Balaguer, Mark. Platonism and Anti-Platonism in Mathematics. New York: Oxford University Press, 1998.

Barrett, Robert, and Gibson, Roger, ed. Perspectives on Quine. Cambridge, Mass.: Blackwell, 1990.

Benacerraf Paul. "Mathematical Truth." In Philosophy of Mathematics, edited by Paul Benacerraf and Hilary Putnam. Cambridge University Press: Cambridge, 2nd edition, 1983.

_____. "What Numbers Could Not Be." In Benacerraf and Putnam (1983).

Boghossian, Paul, and Peacocke, Christopher, ed. New Essays on the A Priori New York: Oxford University Press, 2000.

Bonjour, Lawrence. In Defense of Pure Reason. Cambridge: Cambridge University Press, 1998.

_____. "Toward a Moderate Realism," Philosophical Topics, 23 (1995): 47-78.

_____. The Structure of Empirical Knowledge. Cambridge, Massachusetts: Harvard University Press, 1985.

Burge, Tyler. "Frege on Apriority." In (Boghossian and Peacocke, 2000): 11-42.

_____. "Computer Proof, A Priori Knowledge, and Other Minds." In Philosophical Perspectives, 12, Language, Mind and Ontology, Boston, Mass.: Blackwell, 1998: 1-37.

_____. "Content Preservation," Philosophical Review, V0l. 102, No.4 (October 1993): 457-488.

Casullo, Albert. "A Priori Knowledge Appraised." In A Priori Knowledge, edited by

Abert Casullo. The International Research Library of Philosophy, Aldershot: Dartmouth Publishing Company, 1999.

_____. "Causality, Reliabilism, and Mathematical Knowledge," Philosophy and Phenomenological Research (September 1992): 557-84.

Chihara, Charles. "A Gödelian Thesis Regarding Mathematical Objects: Do They Exist? And Can We Perceive Them?" The Philosophical Review, 91(1982): 211-27.

Clark, Chalmers. "The Art of Science: Quine and the Speculative Reach of Philosophy in Natural Science," Dialectica, Vol. 2, No.4 (1988): 275-290.

Conee, Earl, and Feldman, Richard. "The Generality Problem for Reliabilism." In The Theory of Knowledge, 2nd edition, edited by Louis Pojman. CA:Wadsworth, 1999: 343-357.

Feldman, Richard. "Reliability and Justification," The Monist, (1985): 159-73.

Feldman, Richard, and Conee, Earl. "Evidentialism," Philosophical Studies, 48 (1985): 15-34.

Field, Hartry. "Realism, Mathematics and Modality." In Realism, Mathematics and Modality, Oxford: Blackwell, 1989.

Fodor, Jerry. Modularity of Mind: An Essay on Faculty Psychology. Cambridge, Mass.: MIT Press, 1983.

Foley, Richard. "How Should Future Opinion Affect Current Opinion?" Philosophy and Phenomenological Research (1993): 747-766.

Gettier, Edmund. "Is Justified True Belief Knowledge?" Analysis, Vol. 23 (1963): 121-123.

Gibson, Roger. Enlightened Empiricism. Tampa: University of South Florida Press, 1988.

_____. The Philosophy of W.V.O. Quine, An Expository Essay. Tampa: University of South Florida Press, 1982.

Gödel, Kurt, "What Is Cantor's Continuum Problem?" In (Benacerraf and Putnam 1983).

Goldman, Alvin. Epistemology and Cognition. Cambridge, Ma.: Harvard University Press 1986.

_____."What is Justified Belief?" In Justification and Knowledge, edited by G. Pappas Dordrecht, Reidel (1979): 1-23.

_____. "Discrimination and Perceptual Knowledge," Journal of Philosophy 73 (1976): 771-791.

_____. "A Causal Theory of Knowing," Journal of Philosophy 64 (1967): 357-372.

Grice, H.P., and Strawson, P.F. "In Defense of a Dogma." In Readings in the Philosophy of Language edited by J. Rosenberg, and C. Travis. Englewood Cliffs, N.J.: Prentice Hall, 1971.

Haack, Susan. Evidence and Inquiry: Towards Reconstruction in Epistemology; Cambridge, Massachusetts: Blackwell, 1993.

Hahn, L., and Schilpp, P. The Philosophy of W.V. Quine; La Salle, Ill.: Open Court, 1986.

Hale, Bob. Abstract Objects; Oxford: Blackwell, 1988.

Hale, Bob and Wright, Crispin. "A Reductio Ad Surdum? Field on the Contingency of Mathematical Objects," Mind 103:410 (April 1994): 169-183.

Hardy, G.H. A Mathematician's Apology; London: Cambridge University Press, 1940.

Hart, W.D. "Access and Inference." In The Philosophy of Mathematics, edited by W.D. Hart. New York, Oxford University Press, 1996: 60 – 61.

Hart, W.H. "Review of Mark Steiner's Mathematical Knowledge," Journal of Philosophy 74 (1977).

Horowitz, Tamara. "A Priori Truth," The Journal of Philosophy, Volume 82, Issue 5 (May, 1985): 225 – 239.

Kanigel, Robert. The Man Who Knew Infinity. Scribners (Macdonald & Co.): London, 1992.

Katz, Jerrold. "From the Philosophy of Mathematics to the Philosophy of Philosophy," Journal of Philosophy, forthcoming.

______. Realistic Rationalism. Cambridge, Mass.: Bradford MIT, 1998.

______. "Skepticism About Numbers and Indeterminacy Arguments." In Benacerraf and His Critics, edited by Adam Morton and Stephen Stich. Oxford:

Basil Blackwell, 1996: 119-39.

_____. The Metaphysics of Meaning. Cambridge, Mass.: Bradford MIT, 1990.

_____. "Common Sense in Semantics," Notre Dame Journal of Formal Logic (April 1982): 174-218.

Kim, Jaegwon. "What is 'Naturalized Epistemology'?" In The Theory of Knowledge, edited by Louis Pojman. CA: Wadsworth, 1993.

Kitcher, Philip. "A Priori Knowledge Revisited." In (Boghossian and Peacocke, 2000).

_____. The Nature of Mathematical Knowledge. New York, Oxford, 1983.

Kornblith, Hilary. "What is Naturalistic Epistemology?" In Naturalizing Epistemology, edited by Hilary Kornblith. Cambridge, Ma.: MIT, 1987.

Kripke, Saul. Naming and Necessity. Cambridge: Harvard University Press, 1980.

Levin, Margarita. "On Tymoczko's Argument For Mathematical Empiricism," Philosophical Studies 39 (1981): 79-86.

Levin, Michael. "You Can Always Count On Reliabilism," Philosophy and Phenomenological Research, Vol. LVII, No. 3 (1997): 1-11.

_____. "Reliabilism and Induction," Synthese 97 (1993): 297-334.

Levin, Michael, and Adler, Jonathan. "Is The Generality Problem Too General?" Philosophy and Phenomenological Research, LXV (2002): 87-97.

Maddy, Penelope. Realism in Mathematics. Oxford: Oxford University Press, 1990.

_____. "Perception and Mathematical Intuition," Philosophical Review 89 (1980): 163-96.

Marcus, Ruth Barcan. "A Backwards Look at Quine's Animadversions on Necessity." In (Barrett and Gibson 1990) 230-43.

McEvoy, Mark. "Naturalized Epistemology, Normativity, and the Argument Against the A Priori." Essays in Philosophy, 2002 <www.humboldt.edu/~essays>.

Nozick, Robert. Philosophical Explanations. Cambridge, Massachusetts: Belknap, Harvard University Press, 1981.

Orenstein, Alex. Willlard Van Orman Quine. Boston: Twayne, 1977.

Peacocke, Chris. "Explaining the A Priori: The Program of Moderate Rationalism." In (Boghossian and Peacocke, 2000): 255-85.

Plantinga, Alvin. Warrant: The Current Debate. New York: Oxford University Press, 1993.

______. Warrant and Proper Function, New York, Oxford University Press, 1993.

Pollock, John. "Epistemic Norms," Synthese 71 (1987): 61-95.

_____. Knowledge and Justification. New Jersey: Princeton, 1974.

Putnam, Hilary. "'Two Dogmas' Revisited," in Realism and Reason, Philosophical Papers, Volume 3. Cambridge: Cambridge University Press, 1983.

Quine, W.V.O. "Comment on Lauener." In Barrett and Gibson (1990).

_____. "Comment on Quinton." In (Barrett and Gibson, 1990).

_____. "Reply to Vuillemin." In The Philosophy of W.V. Quine, edited by L. Hahn and P. Schilpp. La Salle, Ill.: Open Court, 1986.

_____. "Reply to White." In (Hahn and Schilpp, 1986).

_____. "Reply to Parsons." In (Hahn and Schilpp, 1986).

_____. "Review of Parson's Mathematics in Philosophy," Journal of Philosophy (1983): 783-94

_____. Theories and Things. Cambridge, Mass.: Harvard University Press, 1981.

_____. "Empirical Content." In (Quine, 1981): 24-30.

_____. "Five Milestones of Empiricism." In (Quine, 1981): 67-72.

_____. "Success and the Limits of Mathematization." In (Quine, 1981): 148-155.

_____. "The Nature of Natural Knowledge." In Mind and Language, edited by S. Guttenplan. Oxford: Clarendon, 1975.

_____. "Epistemology Naturalized." In Ontological Relativity and other Essays. New York; Columbia University Press, 1969.

_____. "Two Dogmas of Empiricism." In From a Logical Point of View (2cd. ed.).

Cambridge, Mass: Harvard University Press, 1961.

_____. Word and Object. Cambridge, Mass.: MIT Press, 1960.

Resnik, Michael. Mathematics as a Science of Patterns. New York: Oxford University Press, 1997.

Sosa, Ernest. "Rational Intuition: Bealer On its Nature and Epistemic Status," Philosophical Studies 81 (1996): 151-162.

Steiner, Mark. Mathematical Knowledge. Ithaca, NY: Cornell University Press 1975.

_____. "Platonism and the Causal Theory of Knowledge," Journal of Philosophy 70 (1973): 57-66.

Stroud, Barry. The Significance of Philosophical Skepticism. Oxford: Oxford University Press, 1984.

Tymoczko, Thomas. "The Four-Color Problem and its Philosophical Significance," Journal of Philosophy, Vol. LXXVI, No.2 (February 1979): 57-82.

Van Fraassen, Bas. "Belief and the Will." Journal of Philosophy (1984): 235-56.

Warner, Richard. "Why Is Logic A Priori?" The Monist 72 (1989): 40-51.

White, Morton. "The Analytic and the Synthetic, an Untenable Dualism." In John Dewey: Philosopher of Science and Freedom, edited by Sydney Hook. New York: Dial, 1950.